arXiv:1211.5620v1 [stat.ME] 23 Nov 2012

Pair-copula Bayesian networks

Alexander Bauer*† Claudia Czado*

Abstract. Pair-copula Bayesian networks (PCBNs) are a novel class of multivariate statistical models, which combine the distributional flexibility of pair-copula constructions (PCCs) with the parsimony of conditional independence models associated with directed acyclic graphs (DAG). We are first to provide generic algorithms for random sampling and likelihood inference in arbitrary PCBNs as well as for selecting orderings of the parents of the vertices in the underlying graphs. Model selection of the DAG is facilitated using a version of the well-known PC algorithm which is based on a novel test for conditional independence of random variables tailored to the PCC framework. A simulation study shows the PC algorithm's high aptitude for structure estimation in non-Gaussian PCBNs. The proposed methods are finally applied to modelling financial return data.

Key words: Conditional independence test; copulas; directed acyclic graphs; graphical models; likelihood inference; PC algorithm; regular vines; structure estimation.

1. Introduction

Graphical models provide a powerful tool in multivariate statistical analysis aimed at modelling the conditional independence structure of a family of random variables. The conditional independence restrictions observed by a graphical model can be conveniently summarised in a graph whose vertices represent the variables and whose edges indicate interrelations between these variables, see Lauritzen (1996). We are particularly interested in the graphical models known as Bayesian networks, whose Markov properties can be represented by a directed acyclic graph (DAG). Areas of applications for these Bayesian networks range from artificial intelligence, decision support systems, and engineering to genetics, geology, medicine, and finance, see Pourret et al. (2008). Despite the broad scope of applicability, however, graphical

*Department of Mathematics, Technische Universität München, Boltzmannstr. 3, 85748 Garching, Germany. E-mail: alexander.bauer@tum.de, cczado@ma.tum.de.
†Corresponding author.

modelling of continuous random variables has mainly been limited to the multivariate normal distribution. Accordingly, available structure estimation algorithms for the DAG underlying a Bayesian network are mainly confined to discrete or Gaussian models. We address both the problems of constructing Bayesian networks with non-Gaussian continuous joint distributions, and of estimating the Markov structure underlying such a non-Gaussian Bayesian network.

Our solution to the first problem of deriving non-Gaussian distributions with pre-specified conditional independence properties is based on so-called pair-copula constructions (PCCs). By iterated application of Sklar's theorem on copulas (Sklar, 1959), Kurowicka and Cooke (2005) and Bauer et al. (2012) have shown that every continuous multivariate distribution associated with a DAG can be decomposed into a family of bivariate, potentially conditional distributions, which correspond to the edges of the underlying graph. An explicit representation of the respective probability density function (pdf) was, however, only derived in examples. We provide a novel algorithm for evaluating the pdf of an arbitrary Bayesian network PCC.

The flexibility of these pair-copula Bayesian networks (PCBNs) allows for the capturing of a wide range of distributional features to be modelled such as heavy-tailedness, tail dependence, and non-linear, asymmetric dependence. Further investigations on PCBNs include Hanea et al. (2006, 2010) and Hanea and Kurowicka (2008). While these authors concentrate on non-parametric statistical inference and elicited expert knowledge, we focus attention to parametric likelihood inference and data-driven structure estimation. We also provide routines for copula selection and enumeration of the parents of the vertices of the underlying DAG.

When expert knowledge on the underlying Markov structure is unavailable, data-driven structure estimation algorithms are frequently used. Two approaches are predominantly found in the literature: the constraint-based and the score-and-search-based approach (Koller and Friedman, 2009, Chapter 18). In the former, the DAG is inferred from a series of conditional independence tests, while in the latter, the DAG is found by optimising a given scoring function. We concentrate on the popular constraint-based PC algorithm by Spirtes and Glymour (1991), and demonstrate its aptitude for structure estimation in non-Gaussian PCBNs in an extensive simulation study. In particular, we introduce a novel test for conditional independence of continuous random variables which is based on the closely related regular-vine copula models (Bedford and Cooke, 2001, 2002), and which is of interest on its own merits. This novel test will prove to outperform a standard test for zero partial correlation used in the Gaussian setting.

With their focus on conditional independence, PCBNs are generally more parsimonious than regular-vine copula models. Another copula decomposition of a joint distribution associated with a DAG which uses generally higher-variate copulas—and therefore lacks the flexibility of the pair-copula approach—was investigated by Elidan (2010, 2012).

The paper is organised as follows. In Section 2, we give a short review of Bayesian networks, followed by a review of vine copula models in Section 3. In Section 4, we provide an algorithm for evaluating the pdf of a PCC associated with a DAG as well as routines for simulation, model selection, and likelihood inference in PCBNs. We review the PC algorithm in Section 5 and introduce a novel test for conditional independence of continuous random variables. The PC algorithm's aptitude for structure estimation in non-Gaussian PCBNs is explored in a simulation study in Section 6. Section 7 presents an application of PCBNs to financial return data, and the paper concludes with a brief discussion in Section 8. The paper is designed to be self-contained and to unify the various non-standard notations on Bayesian networks found in the literature.

2. Bayesian networks

We begin by fixing some graph theoretical terminology. Let $V \neq \emptyset$ be a finite set and let $E \subseteq \mathcal{E} := \{(v,w) \in V \times V \,|\, v \neq w\}$. Then $\mathcal{G} = (V,E)$ denotes a *graph* with *vertex set* V and *edge set* E. We say that $\mathcal{G}$ contains the *undirected edge* $v - w$ if $(v,w) \in E$ and $(w,v) \in E$. Similarly, we say that V contains the *directed edge* $v \to w$ if $(v,w) \in E$ but $(w,v) \notin E$. A graph containing only undirected edges is called an *undirected graph* (UG). If $E \equiv \mathcal{E}$, we call $\mathcal{G}$ the *complete UG* on V. A graph containing only directed edges is called a *directed graph*. By replacing all directed edges of $\mathcal{G}$ with undirected edges, we obtain the *skeleton* $\mathcal{G}^s$ of $\mathcal{G}$. We write $v \multimap w$ whenever $(v,w) \in E$, that is $\mathcal{G}$ contains either the directed edge $v \to w$ or the undirected edge $v - w$. A sequence of distinct vertices $v_1, \ldots, v_k \in V$, $k \geq 2$, is called a *path* from v_1 to v_k if $\mathcal{G}$ contains $v_i \multimap v_{i+1}$ for all $i \in \{1, \ldots, k-1\}$. A path from v_1 to v_k is called *directed* if at least one of the connecting edges is directed. We call a path from v_1 to v_k a *cycle* if $v_1 = v_k$. In particular, we call a directed path from v_1 to v_k a *directed cycle* if $v_1 = v_k$. A graph without directed cycles is called a *chain graph* (CG). A CG containing only directed edges is known as a *directed acyclic graph* (DAG). We define the *adjacency set* of a vertex $v \in V$ as $\mathrm{ad}(v) := \{w \in V \,|\, (v,w) \in E \text{ or } (w,v) \in E\}$. If $w \notin \mathrm{ad}(v)$, we say that v and w are *non-adjacent*. A triple of vertices (u,v,w) is called a *v-structure* if $\mathcal{G}$ contains $u \to v \leftarrow w$ and if u and w are non-adjacent.

Now let $\mathcal{G}$ be a DAG. The *moral graph* $\mathcal{G}^m$ of $\mathcal{G}$ is defined as the skeleton of the graph obtained from $\mathcal{G}$ by introducing an undirected edge $u - w$ whenever $\mathcal{G}$ contains a v-structure (u,v,w) for $u,v,w \in V$. Since all edges of $\mathcal{G}$ are directed, we can speak of paths instead of directed paths. For $v \in V$, we call $\mathrm{pa}(v) := \{w \in V \,|\, \mathcal{G} \text{ contains } w \to v\}$ the *parents* of v, $\mathrm{an}(v) := \{w \in V \,|\, \mathcal{G} \text{ contains a path from } w \text{ to } v\}$ the *ancestors* of v, $\mathrm{de}(v) := \{w \in V \,|\, \mathcal{G} \text{ contains a path from } v \text{ to } w\}$ the *descendants* of v, and $\mathrm{nd}(v) := V \setminus (\{v\} \cup \mathrm{de}(v))$ the *non-descendants* of v. A set $I \subseteq V$ is called *ancestral* if $\mathrm{pa}(v) \subseteq I$ for all $v \in I$. The smallest ancestral set containing I is denoted by $\mathrm{An}(I)$. As is readily verified, $\mathrm{An}(I) = I \cup \bigcup_{v \in I} \mathrm{an}(v)$. The graph

$\mathcal{G}_I = \big(I, E \cap (I \times I)\big)$ is called the *subgraph* of $\mathcal{G}$ induced by I. A bijection $v_\bullet \colon \big\{1, \dots, |V|\big\} \to V$, $i \mapsto v_i$, satisfying $i < j$ whenever $\mathcal{G}$ contains $v_i \to v_j$ for some $i, j \leq |V|$ is called a *well-ordering* of $\mathcal{G}$. Note that in a well-ordered DAG the set $\{v_1, \dots, v_k\}$ is ancestral for all $k \leq |V|$.

Finally, let $\mathcal{G}$ be a UG and let $I, J, K \subseteq V$ be pairwise disjoint. A path from I to J is a path from a vertex $v \in I$ to a vertex $w \in J$. We say that K *separates* I from J in $\mathcal{G}$, and write $I \perp J \mid K\,[\mathcal{G}]$, if every path from I to J contains a vertex in K. In particular, we write $I \perp J \mid \emptyset\,[\mathcal{G}]$, or shortly $I \perp J\,[\mathcal{G}]$, if there exists no path between I and J. We call $\mathcal{G}$ *connected* if for every distinct $v, w \in V$ there is a path from v to w. A connected UG without cycles is a *tree*. If there is a vertex $w \in V$ such that $\mathrm{ad}(w) = V \setminus \{w\}$ and $\mathrm{ad}(v) = \{w\}$ for all $v \in V \setminus \{w\}$, that is all vertices are solely adjacent to w, then $\mathcal{G}$ is called a *star* and w is called its *root vertex*. Note that above terminology is not used consistently throughout the literature.

Markovian probability measures

In graphical probability modelling, graphs are used to represent conditional independence properties of corresponding families of probability measures. Let $\mathcal{D} = (V, E)$ be a DAG on $d := |V|$ vertices and let P be a probability measure on $\mathbb{R}^d$. Moreover, let $\boldsymbol{X}$ be an $\mathbb{R}^d$-valued random variable distributed as P. For $I \subseteq V$, we write $\boldsymbol{X}_I := (X_v)_{v \in I}$ and denote the corresponding I-margin of P by P_I. If $I = \{v\}$ for some $v \in V$, we write X_v and P_v instead of $X_{\{v\}}$ and $P_{\{v\}}$. Furthermore, we write $\boldsymbol{X}_I \perp\!\!\!\perp \boldsymbol{X}_J \mid \boldsymbol{X}_K$ whenever $\boldsymbol{X}_I$ and $\boldsymbol{X}_J$ are *conditionally independent* given $\boldsymbol{X}_K$ for pairwise disjoint sets $I, J, K \subseteq V$. By convention, $\boldsymbol{X}_I \perp\!\!\!\perp \boldsymbol{X}_J \mid \boldsymbol{X}_\emptyset$ is understood as $\boldsymbol{X}_I \perp\!\!\!\perp \boldsymbol{X}_J$. P is said to possess the *local $\mathcal{D}$-Markov* property if

$$X_v \perp\!\!\!\perp \boldsymbol{X}_{\mathrm{nd}(v) \setminus \mathrm{pa}(v)} \mid \boldsymbol{X}_{\mathrm{pa}(v)} \quad \text{for all } v \in V. \tag{2.1}$$

Correspondingly, P is said to possess the *global $\mathcal{D}$-Markov* property if

$$I \perp J \mid K\,[(\mathcal{D}_{\mathrm{An}(I \cup J \cup K)})^m] \;\Rightarrow\; \boldsymbol{X}_I \perp\!\!\!\perp \boldsymbol{X}_J \mid \boldsymbol{X}_K \quad \text{for all pairwise disjoint } I, J, K \subseteq V. \tag{2.2}$$

Equations (2.1) and (2.2) relate (conditional) independence properties of P to graph separation properties of $\mathcal{D}$. Since $\mathrm{ad}(v) \cap \big(\mathrm{nd}(v) \setminus \mathrm{pa}(v)\big) = \emptyset$ for every $v \in V$, it can be easily seen that the conditional independence restrictions obtained from Equation (2.1) correspond to missing edges in $\mathcal{D}$. One can show that P has the local $\mathcal{D}$-Markov property if and only if P has the global $\mathcal{D}$-Markov property, see Lauritzen (1996, p. 51). A probability measure satisfying Equations (2.1) and (2.2) is thus simply called *$\mathcal{D}$-Markovian*. Despite the aforementioned equivalence, the lists of *explicit* conditional independence restrictions obtained from Equations (2.1) and (2.2) may, however, be of different lengths. Note that a $\mathcal{D}$-Markovian probability measure can exhibit further conditional independence properties apart from those represented by $\mathcal{D}$. If, however, P exhibits no conditional independence properties other than those represented by $\mathcal{D}$, then P is

called *faithful* to $\mathcal{D}$. Now let P have Lebesgue-density f. One can show that P is $\mathcal{D}$-Markovian if and only if f has a so-called *$\mathcal{D}$-recursive factorisation*, that is

$$f(\boldsymbol{x}) = \prod_{v \in V} f_{v|\mathrm{pa}(v)}\big(x_v \,\big|\, \boldsymbol{x}_{\mathrm{pa}(v)}\big) \quad \text{for all } \boldsymbol{x} = (x_1, \ldots, x_d) \in \mathbb{R}^d,$$

where $f_{v|\mathrm{pa}(v)}(\,\cdot \mid \boldsymbol{x}_{\mathrm{pa}(v)})$ denotes the conditional probability density function (pdf) of X_v given $\boldsymbol{X}_{\mathrm{pa}(v)} = \boldsymbol{x}_{\mathrm{pa}(v)}$, see again Lauritzen (1996, p. 51). Note that there may be more than one DAG representing the same set of conditional independence restrictions. We call the set of DAGs representing the same conditional independence restrictions as $\mathcal{D}$ the *Markov-equivalence class* of $\mathcal{D}$, and denote it by $[\mathcal{D}]$. Two DAGs $\mathcal{D}_1 = (V, E_1)$ and $\mathcal{D}_2 = (V, E_2)$ are called *Markov equivalent* if $[\mathcal{D}_1] = [\mathcal{D}_2]$. By Verma and Pearl (1991), $\mathcal{D}_1$ and $\mathcal{D}_2$ are Markov equivalent if and only if they have the same skeleton and the same v-structures. The Markov-equivalence class of $\mathcal{D}$ can be represented by a CG, the so-called *essential graph* $\mathcal{D}^e$ associated with $[\mathcal{D}]$, which has the same skeleton as $\mathcal{D}$ and contains a directed edge $v \to w$ if and only if all members of $[\mathcal{D}]$ contain $v \to w$, see Andersson et al. (1997). A DAG in $[\mathcal{D}]$ can be obtained from $\mathcal{D}^e$ by directing all undirected edges of $\mathcal{D}^e$ such that no new v-structures and no directed cycles are introduced. Figure 1 gives an example of a DAG on four vertices together with the essential graph associated with the corresponding Markov-equivalence class.

Figure 1: A DAG $\mathcal{D}$ (left) specifying the conditional independence restrictions $X_1 \perp\!\!\!\perp X_4 \mid \boldsymbol{X}_{23}$ and $X_2 \perp\!\!\!\perp X_3 \mid X_1$, and the essential graph $\mathcal{D}^e$ (right) associated with the corresponding Markov-equivalence class $[\mathcal{D}]$.

Graphical models

A *Bayesian network* or *(directed) graphical model* based on $\mathcal{D}$ is a family of $\mathcal{D}$-Markovian probability measures. A comprehensive introduction to graphical models, and Bayesian networks in particular, is found in Lauritzen (1996) and Cowell et al. (2003), see also Pourret et al. (2008) for examples of applications. For lack of tractable continuous probability measures, statistical modelling with Bayesian networks has mostly been limited to multivariate discrete or normal distributions. Kurowicka and Cooke (2005) therefore used copulas to derive a rich and tractable class of continuous Bayesian networks, which we will investigate in Section 4.

3. Vine copula models

A d-variate *copula*, $d \in \mathbb{N}$, is a cumulative distribution function (cdf) on $[0,1]^d$ such that all univariate marginals are uniform on the interval $[0,1]$. By *Sklar's theorem* (Sklar, 1959), every cdf F on $\mathbb{R}^d$ with marginals $F_1, \dots, F_d$ can be written as

$$F(\boldsymbol{x}) = C\big(F_1(x_1), \dots, F_d(x_d)\big), \quad \boldsymbol{x} = (x_1, \dots, x_d) \in \mathbb{R}^d,$$

for some suitable copula C. If F is absolutely continuous and $F_1, \dots, F_d$ are strictly increasing, a similar relationship holds for the pdf f of F, namely

$$f(\boldsymbol{x}) = c\big(F_1(x_1), \dots, F_d(x_d)\big) \prod_{i=1}^{d} f_i(x_i), \quad \boldsymbol{x} = (x_1, \dots, x_d) \in \mathbb{R}^d,$$

where the *copula density* c is uniquely determined. A comprehensive introduction to copulas is found in Joe (1997) and Nelsen (2006).

3.1. Pair-copula constructions and regular vines

While in recent years a vast catalogue of bivariate copula families (also known as *pair-copula* families) has accumulated in the literature, many of these bivariate families have no straightforward multivariate extension. Based on Joe (1996), Bedford and Cooke (2001, 2002) introduced a rich and flexible class of multivariate copulas that uses bivariate (conditional) copulas as building blocks only. The corresponding decomposition of a multivariate copula into bivariate copulas is called a *pair-copula construction* (PCC). The most widely researched copulas arising from PCCs are the *vine copulas*. These vine copulas admit a graphical representation called a *regular vine* (R-vine), which essentially consists of a sequence of trees, each edge of which is associated with a certain pair copula in the corresponding PCC. More precisely, let $V \neq \emptyset$ be a finite set and let $d := |V|$. An R-vine on V is a sequence $\mathcal{V} := (T_1, \dots, T_{d-1})$ of trees $T_1 = (V_1, E_1), \dots, T_{d-1} = (V_{d-1}, E_{d-1})$ such that $V_1 = V$ and $V_i = E_{i-1}$ for $i \geq 2$, that is the vertices of tree T_i are the edges of tree T_{i-1}. We here represent an edge $v - w$ in tree T_i, $i \in \{1, \dots, d-1\}$, by the doubleton $\{v, w\}$ instead of by the pairs (v, w) and (w, v), that is $E_i \subseteq \big\{\{v, w\} \,\big|\, v \neq w \in V_i\big\}$. Moreover, every tree T_i, $i \geq 2$, of $\mathcal{V}$ has to satisfy a *proximity condition* requiring that $|v \triangle w| = 2$ for every edge $\{v, w\} \in E_i$, where $u \triangle v = (u \cup v) \setminus (u \cap v)$. Two vertices in tree T_i, $i \geq 2$, can hence only be adjacent if the corresponding edges in tree T_{i-1} share a common vertex. Last, every edge $\{v, w\} \in E := E_1 \cup \dots \cup E_{d-1}$ carries a label $v \triangle w \,|\, v \cap w$ representing the (conditional) pair copula $C_{v \triangle w | v \cap w}$, where $v \triangle w \,|\, \emptyset$ is conveniently replaced by $v \triangle w$. Instead of $C_{v \triangle w | v \cap w}$ we also write $C_{v_\triangle, w_\triangle | v \cap w}$, where $v_\triangle := v \setminus (v \cap w)$ and $w_\triangle := w \setminus (v \cap w)$. The pdf f of a d-variate probability measure with univariate marginals F_v,

$v \in V$, and copula $C_{\mathcal{V}}$ corresponding to $\mathcal{V}$ then takes the form

$$f(\boldsymbol{x}) = \prod_{\{v,w\}\in E} c_{v_\triangle, w_\triangle | v\cap w}\big(F_{v_\triangle|v\cap w}(x_{v_\triangle} \,|\, \boldsymbol{x}_{v\cap w}), F_{w_\triangle|v\cap w}(x_{w_\triangle} \,|\, \boldsymbol{x}_{v\cap w}) \,\big|\, \boldsymbol{x}_{v\cap w}\big) \prod_{v\in V} f_v(x_v), \tag{3.1}$$

where $\boldsymbol{x} = (x_v)_{v\in V} \in \mathbb{R}^d$. Note that—similar to DAGs—the vertices in the first tree of $\mathcal{V}$ represent the univariate margins of $C_{\mathcal{V}}$. In contrast to DAGs, however, $\mathcal{V}$ does not have an interpretation in terms of Markov properties of $C_{\mathcal{V}}$. An example of an R-vine representing a five-variate vine copula is given in Figure 2.

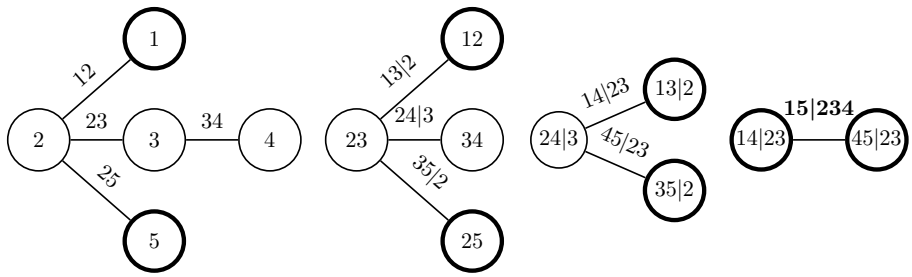

Figure 2: An R-vine specifying the pair copulas C_{12}, C_{23}, C_{25}, C_{34}, $C_{13|2}$, $C_{24|3}$, $C_{35|2}$, $C_{14|23}$, $C_{45|23}$, and $C_{15|234}$ (edge labels). Boundaries of vertices including either 1 or 5 appear in bold, see Section 5.

For every $K \subseteq V$ and $v \in K$, we define $K_{-v} := K \setminus \{v\}$. The conditional cdfs in Equation (3.1) can be evaluated tree-by-tree using a recursive formula derived in Joe (1996), which says that for every $v \in V$, every $K \subseteq V_{-v}$, and an arbitrary $w \in K$

$$F_{v|K}(x_v \,|\, \boldsymbol{x}_K) = \frac{\partial C_{v,w|K_{-w}}\big(F_{v|K_{-w}}(x_v \,|\, \boldsymbol{x}_{K_{-w}}), F_{w|K_{-w}}(x_w \,|\, \boldsymbol{x}_{K_{-w}}) \,\big|\, \boldsymbol{x}_{K_{-w}}\big)}{\partial F_{w|K_{-w}}(x_w \,|\, \boldsymbol{x}_{K_{-w}})}. \tag{3.2}$$

An iterative algorithm for evaluating the pdf in Equation (3.1) under a simplifying assumption of constant conditional copulas introduced below is given in Dißmann et al. (2012). The first partial derivatives of a pair copula $C_{v,w}$ are also known as *h-functions*. We write

$$h_{\underline{v},w}(u_v, u_w) := \frac{\partial C_{v,w}(u_v, u_w)}{\partial u_w} \quad \text{and} \quad h_{v,\underline{w}}(u_v, u_w) := \frac{\partial C_{v,w}(u_v, u_w)}{\partial u_v}, \quad (u_v, u_w) \in [0,1]^2.$$

Many popular pair-copula families exhibit closed-form expressions for these h-functions, see for instance Aas et al. (2009). Note that by Equation (3.2) we have

$$h_{\underline{v},w}\big(F_v(x_v), F_w(x_w)\big) = F_{v|w}(x_v \,|\, x_w) \quad \text{and} \quad h_{v,\underline{w}}\big(F_v(x_v), F_w(x_w)\big) = F_{w|v}(x_w \,|\, x_v),$$

where $(x_v, x_w) \in \mathbb{R}^2$. Hence, we can extend the notion of h-functions to conditional pair copulas,

and express the right hand side of Equation (3.2) by

$$h_{\underline{v},w|K_{-w}}\big(F_{v|K_{-w}}(x_v \,|\, \boldsymbol{x}_{K_{-w}}), F_{w|K_{-w}}(x_w \,|\, \boldsymbol{x}_{K_{-w}}) \,\big|\, \boldsymbol{x}_{K_{-w}}\big).$$

Assume $p := |K| \geq 1$, and write $K = \{w_1, \ldots, w_p\}$ such that $w_i \neq w_j$ for $i \neq j$. We define $K_{-i} := \{w_{i+1}, \ldots, w_p\}$ for every $i \in \{1, \ldots, p\}$. Observing that $f_{v|K}(x_v \,|\, \boldsymbol{x}_K) = \frac{\mathrm{d}}{\mathrm{d}x_v} F_{v|K}(x_v \,|\, \boldsymbol{x}_K)$, we obtain by the chain rule of differentiation

$$f_{v|K}(x_v \,|\, \boldsymbol{x}_K) = f_v(x_v) \prod_{i=1}^{p} c_{v,w_i|K_{-i}}\big(F_{v|K_{-i}}(x_v \,|\, \boldsymbol{x}_{K_{-i}}), F_{w_i|K_{-i}}(x_{w_i} \,|\, \boldsymbol{x}_{K_{-i}}) \,\big|\, \boldsymbol{x}_{K_{-i}}\big). \tag{3.3}$$

3.2. ML estimation and model selection in vine copula models

A *vine copula model* is a family of vine copulas together with families of univariate marginals. Maximum likelihood (ML) estimation in vine copula models was first considered in Aas et al. (2009). The findings therein were, however, restricted to vine copula models represented by C- and D-vines. A *C-vine* is an R-vine whose trees are all stars. Conversely, an R-vine is called a *D-vine* if all vertices in tree T_1 are adjacent to at most two other vertices. ML estimation in vine copula models based on general R-vines was considered in Dißmann et al. (2012).

Let $\mathcal{V}$ be an R-vine on V with edge set E, and let $C_{v_\triangle,w_\triangle|v\cap w}(\,\cdot\,,\cdot\,;\boldsymbol{\theta}_{v_\triangle,w_\triangle|v\cap w})$, $\{v,w\} \in E$, be given (conditional) pair copulas with joint parameter vector $\boldsymbol{\theta} := (\boldsymbol{\theta}_{v_\triangle,w_\triangle|v\cap w})_{\{v,w\}\in E} \in \boldsymbol{\Theta}$. We denote the corresponding vine copula family by $\{C_{\mathcal{V},\boldsymbol{\theta}} \,|\, \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$. Note that we dropped the values $\boldsymbol{x}_{u\cap v}$ of the conditioning variables from the pair copulas $C_{v_\triangle,w_\triangle|v\cap w}$, thus assuming that the corresponding copula family and parameter vector $\boldsymbol{\theta}_{v_\triangle,w_\triangle|v\cap w}$ remain constant for all $\boldsymbol{x}_{u\cap v} \in \mathbb{R}^{|u\cap v|}$. This *simplifying assumption* is made for computational convenience and has become common practice in likelihood inference for vine copula models, see Hobæk Haff et al. (2010) and Acar et al. (2012) for a critical assessment. Furthermore, let $\boldsymbol{u} = \big(\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n\big)$, $n \in \mathbb{N}$, be a realisation of a sample of i.i.d. observations $\boldsymbol{U}^1, \ldots, \boldsymbol{U}^n$ from a random variable $\boldsymbol{U}$ on $[0,1]^d$ with copula family $\{C_{\mathcal{V},\boldsymbol{\theta}} \,|\, \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$ and uniform univariate margins. Equation (3.1) yields the log-likelihood function

$$l(\boldsymbol{\theta};\boldsymbol{u}) = \sum_{k=1}^{n} \sum_{\{v,w\}\in E} \log c_{v_\triangle,w_\triangle|v\cap w}\Big(F_{v_\triangle|v\cap w}\big(u^k_{v_\triangle} \,\big|\, \boldsymbol{u}^k_{v\cap w};\boldsymbol{\theta}\big), F_{w_\triangle|v\cap w}\big(u^k_{w_\triangle} \,\big|\, \boldsymbol{u}^k_{v\cap w};\boldsymbol{\theta}\big);\boldsymbol{\theta}\Big). \tag{3.4}$$

The restriction to uniform univariate margins is made for computational convenience, see below.

ML estimation

Since a joint estimation of the parameters of the univariate marginal distributions and the copula can become computationally demanding in high dimensions, a two-step estimation approach known as the *inference functions for margins* method (Joe and Xu, 1996) is frequently applied.

First, the marginal parameters are estimated and second, given the estimates of the marginal parameters, the copula parameters are inferred. In a similar vein, Genest et al. (1995) proposed a semiparametric approach in which the empirical cdf is used to transform the univariate marginals to uniform $[0, 1]$ distributions before estimating the parameters of the copula model, see Kim et al. (2007) for a comparison. ML estimation of the parameters in Equation (3.4) is frequently performed using a stepwise approach as first described in Aas et al. (2009). In a first step, ML estimates of the parameters of each pair-copula family are computed separately. Due to the recursive structure of the log-likelihood function outlined above, this estimation step is carried out tree-by-tree. We refer to the obtained parameter estimates as *sequential ML estimates*. In a second step, the full log-likelihood function is maximised jointly using the sequential ML estimates as starting values, yielding the so-called *joint ML estimates* $\widehat{\boldsymbol{\theta}}_{v_\triangle, w_\triangle | v \cap w}$, $\{v, w\} \in E$. Large and small sample applications of the stepwise estimation procedure have shown that the sequential ML estimates also provide a good approximation of their joint counterparts, see Hobæk Haff (2012a,b) for consistency results and a simulation study. One might hence consider omitting the second estimation step in a given situation to reduce computational complexity.

Model selection

Model selection for vine copula models comprises an estimation of the R-vine $\mathcal{V}$ and a selection of the pair-copula families for $C_{v_\triangle, w_\triangle | v \cap w}$, $\{v, w\} \in E$. Given $\mathcal{V}$, the latter task of selecting pair-copula families can be performed tree-by-tree, choosing for each edge $\{v, w\} \in E$ the one pair-copula family among a given set of candidate families that optimises a given selection criterion like Akaike's information criterion (AIC) or the Bayesian information criterion (BIC). Dißmann et al. (2012) presented a greedy-type algorithm for the estimation of $\mathcal{V}$, which estimates the trees $T_1, \ldots, T_{d-2}$ sequentially, that is again tree-by-tree. Note that estimating tree T_{d-2} also fixes tree T_{d-1}. Structure estimation for tree $T_i = (V_i, E_i)$, $i \in \{1, \ldots, d-2\}$, is carried out in three steps. In a first step, a weight $\omega_{v,w}$ is assigned to every pair of vertices $v, w \in V_i$ with $|v \triangle w| = 2$. Suitable weights given the data are, for instance, the absolute values of estimates of Kendall's τ, or AIC or BIC values of selected pair-copula families with estimated parameters. In a second step, T_i is set to be a tree on V_i optimising the sum of edge weights $\sum_{\{v,w\} \in E_i} \omega_{v,w}$, where $|v \triangle w| = 2$ for all $\{v, w\} \in E_i$ to ensure the proximity condition. Such an *optimal spanning tree* can be found using the algorithms by Kruskal (1956) or Prim (1957). In a last step, a pair-copula family is assigned to each edge $\{v, w\} \in E_i$, as described above, and an ML estimate of the corresponding parameter(s) is computed. This last step may have already been performed when computing the edge weights $\omega_{v,w}$. Note that due to the greedy nature of the algorithm, the resulting R-vine need not optimise the sum of all edge weights $\sum_{\{v,w\} \in E} \omega_{v,w}$. The search for optimal spanning trees reduces to a search for root vertices when only considering C-vines instead of the more general R-vines, cf. Czado et al. (2012). Since a D-vine is completely determined by tree T_1, only one tree has to be specified when restricting the class of R-vines

to D-vines. Due to the particular structure of D-vines, however, finding tree T_1 by the above method leads to a *travelling salesman problem* (TSP) (Applegate et al., 2007), which is NP-hard. Kurowicka (2011) proposed an alternative structure selection algorithm, in which $\mathcal{V}$ is built in reverse order from tree T_{d-1} to tree T_1 using partial correlation estimates as weights. Bayesian approaches to structure estimation have been considered in Smith et al. (2010), Min and Czado (2011), and Gruber et al. (2012). A more detailed exposition of vine copula models is found in Kurowicka and Joe (2011). Implementations of model selection and ML estimation procedures for vine copula models are available in the R package `VineCopula` (Schepsmeier et al., 2012).

The construction of a d-variate vine copula model requires the specification of $\binom{d}{2}$ pair-copula families, a number growing quadratically in d. The actual number of decisions to make in practical applications may, however, be lower if we happen to discover (conditional) independences in the analysed data. In that case, the corresponding pair copulas are set to be independence copulas. Since above structure estimation algorithm is based on the idea of modelling strongest dependences in the first trees, Brechmann et al. (2012) proposed to set all pair copulas in the later trees to independence copulas, which leads to so-called *truncated R-vines*. Instead of leaving the detection of (conditional) independences to chance, one may, however, consider modelling these independences in the first place to obtain more parsimonious models. Unfortunately, the construction of vine copula models satisfying pre-specified conditional independence restrictions is a hard problem in general. A class of models suited for this task are the Bayesian networks discussed in Section 2. Kurowicka and Cooke (2005) hence joined graphical and copula modelling to introduce PCCs for Bayesian networks, which we will investigate in the next section.

4. Pair-copula Bayesian networks (PCBNs)

Let $\mathcal{D} = (V, E)$ be a DAG, and let P be an absolutely continuous $\mathcal{D}$-Markovian probability measure on $\mathbb{R}^d$, $d := |V|$, with strictly increasing univariate marginal cdfs. Moreover, let $w_v \colon \{1, \dots, |\mathrm{pa}(v)|\} \to \mathrm{pa}(v)$, $i \mapsto w_i := w_v(i)$, be a bijection for every $v \in V$ with $|\mathrm{pa}(v)| \geq 1$. We introduce a total order $<_v$ on $\mathrm{pa}(v)$ for every $v \in V$ such that whenever $|\mathrm{pa}(v)| \geq 1$ we have $w_i <_v w_j$ if and only if $i < j$ for all $i, j \in \{1, \dots, |\mathrm{pa}(v)|\}$. Note that there are $|\mathrm{pa}(v)|!$ permutations of $\mathrm{pa}(v)$ (up to isomorphism). We call $\mathcal{O} := \{<_v \,|\, v \in V\}$ a set of *parent orderings* for $\mathcal{D}$. For every $v \in V$ and $w \in \mathrm{pa}(v)$, we set

$$\mathrm{pa}(v; w) := \{u \in \mathrm{pa}(v) \,|\, u <_v w\} = \{w_i \in \mathrm{pa}(v) \,|\, i < w_v^{-1}(w)\}.$$

By Sklar's theorem, we know that the cdf of P can be uniquely decomposed into the univariate marginals $F_1, \dots, F_d$ and a copula C. Bauer et al. (2012) have shown that C can be further decomposed into the (conditional) pair copulas $C_{v,w|\mathrm{pa}(v;w)}$, $v \in V$, $w \in \mathrm{pa}(v)$, which yields a PCC for C in which each (conditional) pair copula corresponds to exactly one edge $w \to v$ in

$\mathcal{D}$. The pdf f of P can hence be written as

$$f(\boldsymbol{x}) = \prod_{v \in V} f_v(x_v) \prod_{w \in \mathrm{pa}(v)} c_{v,w|\mathrm{pa}(v;w)}\big(F_{v|\mathrm{pa}(v;w)}(x_v \,|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}), F_{w|\mathrm{pa}(v;w)}(x_w \,|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}) \,\big|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}\big), \tag{4.1}$$

where $\boldsymbol{x} = (x_v)_{v \in V} \in \mathbb{R}^d$. As an example consider the DAG in Figure 1 with ordering $2 <_4 3$ of $\mathrm{pa}(4) = \{2, 3\}$. Equation (4.1) yields

$$\begin{aligned} f(\boldsymbol{x}) = {} & f_1(x_1) \cdots f_4(x_4) \cdot c_{21}\big(F_2(x_2), F_1(x_1)\big) \cdot c_{31}\big(F_3(x_3), F_1(x_1)\big) \cdot c_{42}\big(F_4(x_4), F_2(x_2)\big) \\ & \cdot c_{43|2}\big(F_{4|2}(x_4 \,|\, x_2), F_{3|2}(x_3 \,|\, x_2) \,\big|\, x_2\big), \quad \boldsymbol{x} = (x_1, \ldots, x_4) \in \mathbb{R}^4, \end{aligned}$$

where $F_{4|2}(x_4 \,|\, x_2) = h_{\underline{4}2}\big(F_4(x_4), F_2(x_2)\big)$ by Equation (3.2) and

$$F_{3|2}(x_3 \,|\, x_2) = \int_0^1 c_{21}\big(F_2(x_2), u_1\big)\, h_{\underline{3}1}\big(F_3(x_3), u_1\big)\, \mathrm{d}u_1,$$

see Bauer et al. (2012) for details. If we instead choose the ordering $3 <_4 2$ for $\mathrm{pa}(4) = \{2, 3\}$, we obtain the same decomposition as above with the roles of vertices 2 and 3 interchanged. Due to the appearing integral, the pair-copula decomposition in the example cannot be represented by an R-vine. There are, however, DAG PCCs representable by R-vines, see for instance Bauer et al. (2012) for a four-variate DAG PCC which coincides with a D-vine PCC.

4.1. Evaluating conditional cdfs in PCBNs

Similar to vine copulas, the challenge in Equation (4.1) lies in the evaluation of the conditional cdfs. Assume without loss of generality that $\mathcal{D}$ is well-ordered. Let $v \in V$ and let $J \subseteq V \setminus \{v\}$ be non-empty. We will now derive a pair-copula decomposition for the conditional cdf $F_{v|J}(\,\cdot \,|\, \boldsymbol{x}_J)$. We begin by exploiting the (conditional) independence restrictions represented by $\mathcal{D}$. To this end, consider the moral graph $\mathcal{G} := (\mathcal{D}_{\mathrm{An}(\{v\} \cup J)})^m$. If $\{v\} \perp I \mid (J \setminus I)\, [\mathcal{G}]$ for some non-empty $I \subseteq J$, then the global $\mathcal{D}$-Markov property in Equation (2.2) yields with $K := J \setminus I$

$$f_{v|J}(x_v \,|\, \boldsymbol{x}_J) = \frac{f_{\{v\} \cup J}(\boldsymbol{x}_{\{v\} \cup J})}{f_J(\boldsymbol{x}_J)} = \frac{f_{v|K}(x_v \,|\, \boldsymbol{x}_K)\, f_{I|K}(\boldsymbol{x}_I \,|\, \boldsymbol{x}_K)\, f_K(\boldsymbol{x}_K)}{f_{I|K}(\boldsymbol{x}_I \,|\, \boldsymbol{x}_K)\, f_K(\boldsymbol{x}_K)} = f_{v|K}(x_v \,|\, \boldsymbol{x}_K),$$

where by convention $f_{W|\emptyset}(\boldsymbol{x}_W \,|\, \boldsymbol{x}_\emptyset) := f_W(\boldsymbol{x}_W)$ for every $W \subseteq V$, and $f_\emptyset(\boldsymbol{x}_\emptyset) := 1$. Thus, $F_{v|J}(\,\cdot \,|\, \boldsymbol{x}_J) = F_{v|K}(\,\cdot \,|\, \boldsymbol{x}_K)$, and we can continue with the conditioning set K. The case $K = \emptyset$ is trivial. Assume $K \neq \emptyset$. Observing that

$$F_{v|K}(y \,|\, \boldsymbol{x}_K) = \frac{\int_{-\infty}^{y} f_{\{v\} \cup K}(\boldsymbol{x}_{\{v\} \cup K})\, \mathrm{d}x_v}{f_K(\boldsymbol{x}_K)}, \tag{4.2}$$

we next need to find pair-copula decompositions for $f_{\{v\} \cup K}$ and f_K.

Pair-copula decompositions for marginal pdfs

More generally, let $I \subseteq V$ be non-empty and consider the (marginal) pdf f_I. For every $v \in I$, we set $I_{-v} := I \setminus \{v\}$ and obtain the following lemma.

Lemma 4.1. *Let $\mathcal{D} = (V, E)$ be a well-ordered DAG on $d := |V|$ vertices, and let P be an absolutely continuous $\mathcal{D}$-Markovian probability measure on $\mathbb{R}^d$ with pdf f. Let $I \subseteq V$ be non-empty and let v denote the maximal vertex in I by the well-ordering of $\mathcal{D}$. Moreover, define $S_v := \big\{u \in \mathrm{pa}(v) \,\big|\, \{u\} \perp I_{-v}\,[(\mathcal{D}_{\mathrm{An}(\{u\}\cup I_{-v})})^m]\big\}$ and*

$$W_v := \begin{cases} \emptyset & \text{if } I_{-v} = \emptyset \text{ or } S_v = \mathrm{pa}(v), \\ \{w_1\} \cup \mathrm{pa}(v; w_1) & \text{if } I_{-v} \subseteq \mathrm{pa}(v) \text{ and } I_{-v} \neq \emptyset, \\ \{w_2\} \cup \mathrm{pa}(v; w_2) & \text{else,} \end{cases}$$

where w_1 and w_2 denote the maximal vertices in I_{-v} and $\mathrm{pa}(v) \setminus S_v$, respectively, by a given parent ordering $<_v$. Then for all $\boldsymbol{x}_I = (x_u)_{u \in I} \in R^{|I|}$,

$$f_I(\boldsymbol{x}_I) = \int_{\mathbb{R}^{|W_v \setminus I|}} f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v})\, f_{W_v \cup I_{-v}}(\boldsymbol{x}_{W_v \cup I_{-v}})\, \mathrm{d}\boldsymbol{x}_{W_v \setminus I}. \tag{4.3}$$

Note that by convention, $\int_{\mathbb{R}^0} g(\boldsymbol{x})\, \mathrm{d}\boldsymbol{x}_\emptyset := g(\boldsymbol{x})$ for every integrable function $g\colon \mathbb{R}^k \to \mathbb{R}$, $k \in \mathbb{N}$. Also note that the parent ordering $<_v$ need not concur with the well-ordering of $\mathcal{D}$.

Proof. As can be seen from the definition of W_v, the decomposition of f_I in the lemma's claim depends on the relation between the sets I_{-v} and $\mathrm{pa}(v)$. Assume first that $I_{-v} = \emptyset$. Then $f_I = f_v$ and the claim is trivial.

Next, assume $I_{-v} \neq \emptyset$ but $\mathrm{pa}(v) = \emptyset$. Then $S_v = \emptyset$. Since v is maximal in I by the well-ordering of $\mathcal{D}$, v has no descendants in I_{-v}, and we have $\{v\} \perp I_{-v}\,[(\mathcal{D}_{\mathrm{An}(I)})^m]$. The global $\mathcal{D}$-Markov property thus yields $f_I(\boldsymbol{x}_I) = f_v(x_v)\, f_{I_{-v}}(\boldsymbol{x}_{I_{-v}})$, that is Equation (4.3) for $W_v := \emptyset$.

From now on assume $I_{-v} \neq \emptyset$ and $\mathrm{pa}(v) \neq \emptyset$. The possible relations between I_{-v} and $\mathrm{pa}(v)$ are illustrated in Figure 3. If $I_{-v} \subseteq \mathrm{pa}(v)$ (Figure 3a), we extend I_{-v} to $W_v := \{w_1\} \cup \mathrm{pa}(v; w_1) \supseteq I_{-v}$ and obtain as claimed

$$f_I(\boldsymbol{x}_I) = \int_{\mathbb{R}^{|W_v \setminus I|}} f_{\{v\}\cup W_v}(\boldsymbol{x}_{\{v\}\cup W_v})\, \mathrm{d}\boldsymbol{x}_{W_v \setminus I} = \int_{\mathbb{R}^{|W_v \setminus I|}} f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v})\, f_{W_v}(\boldsymbol{x}_{W_v})\, \mathrm{d}\boldsymbol{x}_{W_v \setminus I}.$$

Note that in case $I_{-v} = \{w_1\} \cup \mathrm{pa}(v; w_1)$, no integration is required since then $W_v \setminus I = \emptyset$.

Next, let $I_{-v} \cap \mathrm{pa}(v) = \emptyset$ (Figure 3b). If $S_v = \mathrm{pa}(v)$, then $\{v\} \perp I_{-v}\,[(\mathcal{D}_{\mathrm{An}(I)})^m]$ since v has no descendants in I_{-v}. Hence, we again have $f_I(\boldsymbol{x}_I) = f_v(x_v)\, f_{I_{-v}}(\boldsymbol{x}_{I_{-v}})$, that is Equation (4.3)

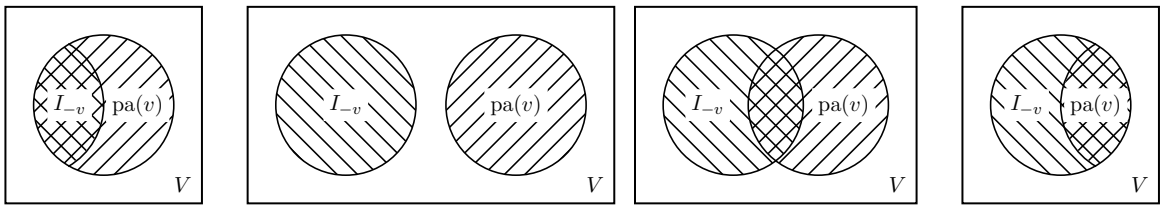

(a) $W_v = \{w_1\} \cup \mathrm{pa}(v; w_1)$ (b) $W_v = \emptyset$ or $W_v = \{w_2\} \cup \mathrm{pa}(v; w_2)$ (c) $W_v = \{w_2\} \cup \mathrm{pa}(v; w_2)$ (d) $W_v = \mathrm{pa}(v)$

Figure 3: Venn diagrams of the sets $I_{-v} \neq \emptyset$ and $\mathrm{pa}(v) \neq \emptyset$, and corresponding definitions of W_v (see lower captions).

for $W_v := \emptyset$. If, however, $S_v \neq \mathrm{pa}(v)$, then $\{v\} \perp I_{-v} \mid \mathrm{pa}(v; w_2)\, [(\mathcal{D}_{\mathrm{An}(\mathrm{pa}(v;w_2)\cup I)})^m]$, and with $W_v := \{w_2\} \cup \mathrm{pa}(v; w_2)$ the global $\mathcal{D}$-Markov property yields

$$f_{I\cup W_v}(\boldsymbol{x}_{I\cup W_v}) = f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v})\, f_{I_{-v}|W_v}(\boldsymbol{x}_{I_{-v}} \,|\, \boldsymbol{x}_{W_v})\, f_{W_v}(\boldsymbol{x}_{W_v}). \tag{4.4}$$

Since $I_{-v} \cap W_v = \emptyset$, we thus get

$$f_I(\boldsymbol{x}_I) = \int_{\mathbb{R}^{|W_v|}} f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v})\, f_{W_v\cup I_{-v}}(\boldsymbol{x}_{W_v\cup I_{-v}})\, \mathrm{d}\boldsymbol{x}_{W_v}.$$

Note that in case $S_v = \emptyset$, we have $W_v = \mathrm{pa}(v)$.

Finally, assume $I_{-v} \cap \mathrm{pa}(v) \neq \emptyset$ such that $I_{-v} \nsubseteq \mathrm{pa}(v)$ (Figure 3c). Similarly to Equation (4.4), we obtain with $W_v := \{w_2\} \cup \mathrm{pa}(v; w_2)$

$$f_{I\cup W_v}(\boldsymbol{x}_{I\cup W_v}) = f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v})\, f_{(I_{-v}\setminus W_v)|W_v}(\boldsymbol{x}_{I_{-v}\setminus W_v} \,|\, \boldsymbol{x}_{W_v})\, f_{W_v}(\boldsymbol{x}_{W_v})$$

by the global $\mathcal{D}$-Markov property, and hence

$$f_I(\boldsymbol{x}_I) = \int_{\mathbb{R}^{|W_v\setminus I|}} f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v})\, f_{W_v\cup I_{-v}}(\boldsymbol{x}_{W_v\cup I_{-v}})\, \mathrm{d}\boldsymbol{x}_{W_v\setminus I}.$$

Note again that in case $S_v = \emptyset$, we have $W_v = \mathrm{pa}(v)$. Also, note that in case $\mathrm{pa}(v) \subseteq I_{-v}$ (Figure 3d), no integration is required. This establishes the claim. □

The set W_v in Lemma 4.1 is either empty or of the form $\{w\} \cup \mathrm{pa}(v; w)$ for some $w \in \mathrm{pa}(v)$. In the latter case, we can express the conditional pdf $f_{v|\{w\}\cup\mathrm{pa}(v;w)}(\,\cdot \,|\, \boldsymbol{x}_{\{w\}\cup\mathrm{pa}(v;w)})$ on the right hand side of Equation (4.3) in terms of the univariate marginals F_u, $u \in V$, and the (conditional) pair copulas $C_{v,u|\,\mathrm{pa}(v;u)}$, $u \in \mathrm{pa}(v)$, as follows.

Lemma 4.2. *Let the notation be as in Lemma 4.1 and let P have strictly increasing univariate marginal cdfs. Let $I_{-v} \neq \emptyset$ and let $S_v \neq \mathrm{pa}(v)$. Then*

$$f_{v|W_v}(x_v \,|\, \boldsymbol{x}_{W_v}) = f_v(x_v) \prod_{w\in W_v} c_{v,w|\mathrm{pa}(v;w)}\big(F_{v|\mathrm{pa}(v;w)}(x_v \,|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}), F_{w|\mathrm{pa}(v;w)}(x_w \,|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}) \,\big|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}\big)$$

for all $x_v \in \mathbb{R}$ *and* $\boldsymbol{x}_{W_v} = (x_w)_{w \in W_v} \in \mathbb{R}^{|W_v|}$.

Proof. Since $I_{-v} \neq \emptyset$ and $S_v \neq \mathrm{pa}(v)$, W_v is non-empty and thus $W_v = \{u\} \cup \mathrm{pa}(v; u)$ for some $u \in \mathrm{pa}(v)$. By Equation (3.3), we can hence write

$$f_{v|W_v}(x_v|\boldsymbol{x}_{W_v}) = f_v(x_v) \prod_{w \in W_v} c_{v,w|\mathrm{pa}(v;w)}\big(F_{v|\mathrm{pa}(v;w)}(x_v|\boldsymbol{x}_{\mathrm{pa}(v;w)}), F_{w|\mathrm{pa}(v;w)}(x_w|\boldsymbol{x}_{\mathrm{pa}(v;w)})\big|\boldsymbol{x}_{\mathrm{pa}(v;w)}\big),$$

and the claim is proven. □

Algorithm 1 Pair-copula decomposition of a (marginal) pdf.

Input Well-ordered DAG $\mathcal{D}$; set of parent orderings $\mathcal{O}$; non-empty vertex set $I \subseteq V$.
Output Factorisation f. % (marginal) pdf $f_I(\boldsymbol{x}_I)$

```
f ← 1;
J ← ∅; % indices of integration variables
while |I| ≥ 1 do
    % Select maximal vertex:
    v ← maximal vertex in I by the well-ordering of D;
    f ← f · f_v(x_v);
    I ← I_{-v};
    % Determine the set W_v:
    W ← ∅;
    S ← {w ∈ pa(v) | {w} ⊥ I [(D_{An({w}∪I)})^m]};
    if I ≠ ∅ and S ≠ pa(v) then
        if I ⊆ pa(v) then
            w ← maximal vertex in I by the parent ordering <_v;
            W ← {w} ∪ pa(v; w);
        else
            w ← maximal vertex in pa(v) \ S by the parent ordering <_v;
            W ← {w} ∪ pa(v; w);
        end if
    end if
    % Introduce corresponding pair copulas and integration variables:
    for w ∈ W do
        f ← f · c_{v,w|pa(v;w)}(F_{v|pa(v;w)}(x_v|x_{pa(v;w)}), F_{w|pa(v;w)}(x_w|x_{pa(v;w)}) | x_{pa(v;w)});
        if w ∉ I then
            I ← I ∪ {w};
            J ← J ∪ {w};
        end if
    end for
end while
f ← ∫_{R^{|J|}} f dx_J;
```

Since all vertices in $W_v \cup I_{-v}$ are smaller than v by the well-ordering of $\mathcal{D}$ and since V is finite, we can inductively apply Lemmas 4.1 and 4.2 to the pdf $f_{W_v \cup I_{-v}}$ in Equation (4.3) until

no unconditional pdfs of dimension higher than one remain. Let J denote the set of vertices corresponding to the integration variables added during this iterative procedure (and including $W_v \setminus I$). Given a set $\mathcal{O}$ of parent orderings for $\mathcal{D}$, Lemma 4.1 yields a set W_u for every $u \in I \cup J$. We have hence established the following theorem.

Theorem 4.3. *Let the notation be as in Lemma 4.2. Then $f_I(\boldsymbol{x}_I)$ takes the form*

$$\int_{\mathbb{R}^{|J|}} \prod_{v \in (I \cup J)} f_v(x_v) \prod_{w \in W_v} c_{v,w|\mathrm{pa}(v;w)}\big(F_{v|\mathrm{pa}(v;w)}(x_v \,|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}), F_{w|\mathrm{pa}(v;w)}(x_w \,|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}) \,\big|\, \boldsymbol{x}_{\mathrm{pa}(v;w)}\big)\, \mathrm{d}\boldsymbol{x}_J$$

for all $\boldsymbol{x}_I = (x_v)_{v \in I} \in R^{|I|}$. □

Note that in the special case $I = V$, Theorem 4.3 yields Equation (4.1). Above procedure for deriving a pair-copula decomposition of f_I as given in Theorem 4.3 is summarised in Algorithm 1.

Example. We consider the well-ordered DAG $\mathcal{D}$ in Figure 4. The edges and parent orderings of $\mathcal{D}$ can be summarised in a matrix $A_{\mathcal{D}} = (a_{ij})_{1 \le i,j \le 7}$ whose elements satisfy $a_{ij} = k$, $k \le |\mathrm{pa}(j)|$, if $\mathcal{D}$ contains the edge $i \to j$ and if i is the k-th smallest parent of j by $<_j$, and $a_{ij} = 0$ otherwise, see Figure 4. For the reader's convenience we will omit function arguments. Equation (4.1) yields

$$\begin{aligned} f = f_1 \cdots f_7 \, & c_{21}(F_2, F_1)\, c_{31}(F_3, F_1)\, c_{42}(F_4, F_2)\, c_{41|2}(F_{4|2}, F_{1|2})\, c_{54}(F_5, F_4)\, c_{53|4}(F_{5|4}, F_{3|4}) \\ & \cdot c_{65}(F_6, F_5)\, c_{64|5}(F_{6|5}, F_{4|5})\, c_{63|54}(F_{6|54}, F_{3|54})\, c_{62|543}(F_{6|543}, F_{2|543})\, c_{75}(F_7, F_5) \\ & \cdot c_{76|5}(F_{7|5}, F_{6|5})\, c_{73|56}(F_{7|56}, F_{3|56}). \end{aligned}$$

We will later derive a pair-copula decomposition for $F_{3|56}$. In preparation, we now use Algorithm 1 to derive pair-copula decompositions for f_{356} and f_{56}.

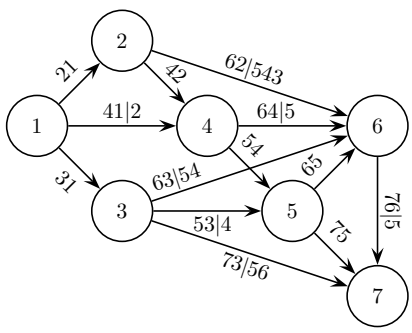

$$\begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \end{array} \begin{array}{c} \begin{array}{ccccccc} 1 & 2 & 3 & 4 & 5 & 6 & 7 \end{array} \\ \begin{bmatrix} 0 & 1 & 1 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 2 & 3 & 3 \\ 0 & 0 & 0 & 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{array}$$

Figure 4: A well-ordered (vertex labels) DAG $\mathcal{D}$ (left) with parent orderings $2 <_4 1$, $4 <_5 3$, $5 <_6 4 <_6 3 <_6 2$, $5 <_7 6 <_7 3$ specifying the pair copulas C_{21}, C_{31}, C_{42}, $C_{41|2}$, C_{54}, $C_{53|4}$, C_{65}, $C_{64|5}$, $C_{63|54}$, $C_{62|542}$, C_{75}, $C_{76|5}$, $C_{73|56}$ (edge labels), and corresponding representation matrix $A_{\mathcal{D}}$ (right).

As a result of applying Algorithm 1 to f_{356} and f_{56}, we obtain

$$\begin{aligned} f_{356} &= \int_{\mathbb{R}^3} f_{6|543}\, f_{5|43}\, f_{4|21}\, f_{3|1}\, f_{2|1}\, f_1 \,\mathrm{d}\boldsymbol{x}_{124} \\ &= \int_{\mathbb{R}^3} f_6\, c_{63|54}(F_{6|54}, F_{3|54})\, c_{64|5}(F_{6|5}, F_{4|5})\, c_{65}(F_6, F_5)\, f_5\, c_{53|4}(F_{5|4}, F_{3|4})\, c_{54}(F_5, F_4) \\ &\quad \cdot f_4\, c_{41|2}(F_{4|2}, F_{1|2})\, c_{42}(F_4, F_2)\, f_3\, c_{31}(F_3, F_1)\, f_2\, c_{21}(F_2, F_1)\, f_1 \,\mathrm{d}\boldsymbol{x}_{124} \end{aligned} \tag{4.5}$$

and $f_{56} = f_{6|5}\, f_5 = f_6\, c_{65}(F_6, F_5)\, f_5$, respectively, see Table 1.

f_{356}	I	v	S	$I_{-v} \subseteq \mathrm{pa}(v)$?	w	W	J
	$\{3,5,6\}$	6	$\emptyset$	✓	3	$\{3,4,5\}$	$\{4\}$
	$\{3,4,5\}$	5	$\emptyset$	✓	3	$\{3,4\}$	$\{4\}$
	$\{3,4\}$	4	$\emptyset$	✗	1	$\{1,2\}$	$\{1,2,4\}$
	$\{1,2,3\}$	3	$\emptyset$	✗	1	$\{1\}$	$\{1,2,4\}$
	$\{1,2\}$	2	$\emptyset$	✓	1	$\{1\}$	$\{1,2,4\}$
	$\{1\}$	1	$\emptyset$	–	–	$\emptyset$	$\{1,2,4\}$
f_{56}	I	v	S	$I_{-v} \subseteq \mathrm{pa}(v)$?	w	W	J
	$\{5,6\}$	6	$\emptyset$	✓	5	$\{5\}$	$\emptyset$
	$\{5\}$	5	$\{3,4\}$	–	–	$\emptyset$	$\emptyset$

Table 1: Vertices and vertex sets obtained during the application of Algorithm 1 to the pdfs f_{356} and f_{56} corresponding to the DAG $\mathcal{D}$ in Figure 4.

When can $\int_{-\infty}^{y} f_{\{v\}\cup K}(\boldsymbol{x}_{\{v\}\cup K})\,\mathrm{d}x_v$ in Equation (4.2) be further simplified?

Let us now return to the conditional cdf in Equation (4.2). Setting $I := \{v\} \cup K$, the numerator on the right hand side of Equation (4.2) takes the form $\int_{-\infty}^{y} f_I(\boldsymbol{x}_I)\,\mathrm{d}x_v$. Decompose $f_I(\boldsymbol{x}_I)$ according to Theorem 4.3, and let J denote the set of vertices corresponding to the newly added integration variables. Clearly, $J \subseteq \mathrm{An}(I) \setminus I$. If the old integration variable x_v does not appear as a conditioning variable in one of the pair copulas $C_{v,w|\mathrm{pa}(v;w)}$, $v \in I \cup J$, $w \in W_v$, in the decomposition of f_I, it may be possible to solve the integral with respect to x_v analytically. More precisely, let $J' \subseteq J$ and let $W' \subseteq W := I_{-v} \cup J'$ be non-empty. Let $k := |W'|$ and write $W' = \{w_1, \ldots, w_k\}$. Moreover, set $W'_{-i} := \{w_1, \ldots, w_{i-1}\}$ for all $i \in \{1, \ldots k\}$. Assume that the pair copula $C_{v,w_k|W'_{-k}}$ is available in the pair-copula decomposition of f, that is $W'_{-k} = \mathrm{pa}(v; w_k)$ or $W'_{-k} = \mathrm{pa}(w_k; v)$, and that (after possible algebraic manipulation) f_I takes the form

$$\int_{\mathbb{R}^{|J'|}} f_v(x_v) \prod_{i=1}^{k} c_{v,w_i|W'_{-i}}\big(F_{v|W'_{-i}}(x_v \mid \boldsymbol{x}_{W'_{-i}}), F_{w_i|W'_{-i}}(x_{w_i} \mid \boldsymbol{x}_{W'_{-i}}) \,\big|\, \boldsymbol{x}_{W'_{-i}}\big)\, f_W(\boldsymbol{x}_W)\,\mathrm{d}\boldsymbol{x}_{J'}. \tag{4.6}$$

Then Fubini's theorem and Equation (3.3) yield that $\int_{-\infty}^{y} f_I(\boldsymbol{x}_I)\,\mathrm{d}x_v$ takes the form

$$\int_{\mathbb{R}^{|J'|}} h_{\underline{v},w_k|W'_{-k}}\big(F_{v|W'_{-k}}(y\,|\,\boldsymbol{x}_{W'_{-k}}), F_{w_k|W'_{-k}}(x_{w_k}\,|\,\boldsymbol{x}_{W'_{-k}})\,\big|\,\boldsymbol{x}_{W'_{-k}}\big)\, f_W(\boldsymbol{x}_W)\,\mathrm{d}\boldsymbol{x}_{J'}, \tag{4.7}$$

where the integral with respect to x_v was replaced by an h-function which, by assumption, is available in the pair-copula decomposition of f. Note that some of the copula pdfs $c_{v,w_i|W'_{-i}}$, $i \in \{1,\dots,k-1\}$, in Equation (4.6) may not correspond to an edge in $\mathcal{D}$, but may instead be given implicitly by an integral over further variables, or may be equal to 1 due to a related Markov property of P, see also the example below. We need to take these special cases into account when checking the applicability of the inverse chain rule algorithmically.

It may sometimes also be useful to substitute $u_w := F_w(x_w)$, that is $\mathrm{d}u_w = f_w(x_w)\,\mathrm{d}x_w$, for all $w \in J$ in the pair-copula decomposition of f_I, and thus to write

$$f_I(\boldsymbol{x}_I) = \int_{[0,1]^{|J|}} c_{I\cup J}\Big(\big(F_w(x_w)\big)_{w\in I}, \boldsymbol{u}_J\Big) \prod_{w\in I} f_w(x_w)\,\mathrm{d}\boldsymbol{u}_J.$$

A similar transformation can be applied to the denominator in Equation (4.2) if integration variables are present.

Example (continued)**.** Consider the integral $\int_{-\infty}^{\cdot} f_{356}\,\mathrm{d}x_3$ associated to the DAG $\mathcal{D}$ in Figure 4, which will later appear when deriving a pair-copula decomposition for $F_{3|56}$. Observing that

$$\begin{aligned}\int_{\mathbb{R}^2} f_4\, c_{41|2}(F_{4|2}, F_{1|2})\, c_{42}(F_4, F_2)\, f_3\, c_{31}(F_3, F_1)\, f_2\, c_{21}(F_2, F_1)\, f_1\,\mathrm{d}\boldsymbol{x}_{12} &= \int_{\mathbb{R}^2} f_{1234}\,\mathrm{d}\boldsymbol{x}_{12} = f_{34}\\ &= f_4\, f_3\, c_{43}(F_4, F_3),\end{aligned}$$

Equation (4.5) yields

$$\begin{aligned}\int_{-\infty}^{\cdot} f_{356}\,\mathrm{d}x_3 = \int_{-\infty}^{\cdot}\int_{\mathbb{R}} & f_3\, c_{63|54}(F_{6|54}, F_{3|54})\, c_{53|4}(F_{5|4}, F_{3|4})\, c_{43}(F_4, F_3)\\ &\cdot f_6\, c_{64|5}(F_{6|5}, F_{4|5})\, c_{65}(F_6, F_5)\, f_5\, c_{54}(F_5, F_4)\, f_4\,\mathrm{d}x_4\,\mathrm{d}x_3.\end{aligned}$$

Note that c_{43} is not available in the pair-copula decomposition of f. Since by Equation (3.3)

$$\int_{-\infty}^{\cdot} f_3\, c_{63|54}(F_{6|54}, F_{3|54})\, c_{53|4}(F_{5|4}, F_{3|4})\, c_{43}(F_4, F_3)\,\mathrm{d}x_3 = h_{6\underline{3}|54}(F_{6|54}, F_{3|54}),$$

we can, however, simplify the integral with respect to x_3, and c_{43} vanishes. We obtain

$$\int_{-\infty}^{\cdot} f_{356}\,\mathrm{d}x_3 = \int_{\mathbb{R}} f_6\, h_{6\underline{3}|54}(F_{6|54}, F_{3|54})\, c_{64|5}(F_{6|5}, F_{4|5})\, c_{65}(F_6, F_5)\, f_5\, c_{54}(F_5, F_4)\, f_4\,\mathrm{d}x_4. \tag{4.8}$$

Pair-copula decompositions for conditional cdfs

Summing up, a pair-copula decomposition for the conditional cdf $F_{v|K}(\,\cdot\,|\,\boldsymbol{x}_K)$ in Equation (4.2) is obtained in three steps. First, we apply Theorem 4.3 to $f_{\{v\}\cup K}$ and f_K. Second, we possibly apply the inverse chain rule to the integral with respect to x_v in the numerator. Last, we cancel common factors like $\prod_{w\in K} f_w(x_w)$ in the numerator and the denominator. The procedure is summarised in Algorithm 2.

Algorithm 2 Pair-copula decomposition of a conditional cdf.

Input Well-ordered DAG $\mathcal{D}$; set of parent orderings $\mathcal{O}$; vertex $v \in V$ (conditioned variable), vertex set $K \subseteq V_{-v}$ (conditioning variables).

Output Factorisation F. % conditional cdf $F_{v|K}(y\,|\,\boldsymbol{x}_K)$

```
% Exploit global D-Markov property:
while ∃ w ∈ K : {v} ⊥ {w} | K_{-w} [(D_{An({v}∪K)})^m] do
    K ← K_{-w};
end while
% Numerator:
n ← Algorithm_1(D, O, {v} ∪ K);
n ← ∫_{-∞}^{y} n dx_v;
simplify n with inverse chain rule for variable x_v if possible;
% Denominator:
d ← Algorithm_1(D, O, K);
% Conditional cdf F_{v|K}(y | x_K):
cancel common factors in n and d;
F ← n/d;
```

As can be seen from Theorem 4.3 and Equation (4.7), the factorisation for $F_{v|K}(\,\cdot\,|\,\boldsymbol{x}_K)$ obtained from Algorithm 2 may contain some new conditional cdfs. This problem can, however, be solved inductively. Let w denote the maximal vertex in $\{v\}\cup K$ by the well-ordering of $\mathcal{D}$. Since Algorithm 2 only adds ancestors of $\{v\}\cup K$ as integration variables, all vertices involved in the new conditional cdfs are smaller than or equal to w by the well-ordering of $\mathcal{D}$. In particular, those conditional cdfs involving w are of the special form $F_{w|\mathrm{pa}(w;u)}(\,\cdot\,|\,\boldsymbol{x}_{\mathrm{pa}(w;u)})$ for some $u \in \mathrm{pa}(w)$, and can by Equation (3.2) iteratively be expressed as

$$F_{v|\mathrm{pa}(w;u)}(x_v\,|\,\boldsymbol{x}_{\mathrm{pa}(w;u)}) = h_{\underline{v},u|\mathrm{pa}(v;u)}\big(F_{v|\mathrm{pa}(v;u)}(x_v\,|\,\boldsymbol{x}_{\mathrm{pa}(v;u)}), F_{u|\mathrm{pa}(v;u)}(x_u\,|\,\boldsymbol{x}_{\mathrm{pa}(v;u)})\,\big|\,\boldsymbol{x}_{\mathrm{pa}(v;u)}\big).$$

Hence, all vertices involved in the algorithmically more demanding new conditional cdfs are strictly smaller than w by the well-ordering of $\mathcal{D}$. Corresponding pair-copula decompositions for the new conditional cdfs can thus be computed inductively by again applying Algorithm 2. Since V is finite, the whole procedure terminates after finitely many steps, and the desired decomposition in terms of only univariate marginals and (conditional) pair copulas is obtained.

Overall, we observed that the problems of deriving pair-copula decompositions for a conditional cdf and a marginal pdf are deeply intertwined and can be solved by alternating iteration. Note that it is sufficient for our purposes to exploit only those conditional independence properties of P which follow directly from graph separation in $\mathcal{D}$ via the global $\mathcal{D}$-Markov property. Once a complete decomposition for f is obtained, the evaluation at $\boldsymbol{x} \in \mathbb{R}^d$ can be performed vertex-by-vertex and parent-by-parent along the well-ordering of $\mathcal{D}$. That is, given $v^* \in V$ and $w^* \in \mathrm{pa}(v^*)$, we first evaluate all terms corresponding to the marginals F_v and the pair copulas $C_{v,w|\mathrm{pa}(v;w)}$ for v smaller than v^* by the well-ordering of $\mathcal{D}$ and $w <_v w^*$ if $w^* \in \mathrm{pa}(v)$, before evaluating the terms corresponding to F_{v^*} and $C_{v^*,w^*|\mathrm{pa}(v^*;w^*)}$.

Example (continued)**.** For the DAG $\mathcal{D}$ in Figure 4, we sketch how to apply Algorithm 2 to obtain a pair-copula decomposition for $F_{3|56}$. Note that $\mathcal{D}$ contains the edges $3 \to 5$ and $3 \to 6$, which is why neither 5 nor 6 can be removed from the conditioning set. We get $F_{3|56} = \int_{-\infty}^{\cdot} \frac{f_{356}}{f_{56}} \,\mathrm{d}x_3$. Applying our previous results for f_{356} and f_{56}, respectively, we further have

$$F_{3|56} = \frac{\int_{\mathbb{R}} f_6 \, h_{6\underline{3}|54}(F_{6|54}, F_{3|54}) \, c_{64|5}(F_{6|5}, F_{4|5}) \, c_{65}(F_6, F_5) \, f_5 \, c_{54}(F_5, F_4) \, f_4 \,\mathrm{d}x_4}{f_6 \, c_{65}(F_6, F_5) \, f_5},$$

see Equation (4.8). Thus, by cancelling common factors, we finally obtain

$$F_{3|56} = \int_{\mathbb{R}} h_{6\underline{3}|54}(F_{6|54}, F_{3|54}) \, c_{64|5}(F_{6|5}, F_{4|5}) \, c_{54}(F_5, F_4) \, f_4 \,\mathrm{d}x_4.$$

4.2. Simulation, ML estimation, and model selection in PCBNs

Given (conditional) pair copulas $C_{v,w|\mathrm{pa}(v;w)}(\,\cdot\,,\cdot\,;\boldsymbol{\theta}_{v,w|\mathrm{pa}(v;w)})$, $v \in V$, $w \in \mathrm{pa}(v)$, with joint parameter vector $\boldsymbol{\theta} := (\boldsymbol{\theta}_{v,w|\mathrm{pa}(v;w)})_{v \in V, w \in \mathrm{pa}(v)} \in \boldsymbol{\Theta}$, above construction yields a d-variate copula model, which we will denote by $\{C_{\mathcal{D},\mathcal{O},\boldsymbol{\theta}} \,|\, \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$. Note that for computational convenience, we again make the *simplifying assumption* of constant conditional copulas described in Section 3.2. Together with families of univariate marginals, $\{C_{\mathcal{D},\mathcal{O},\boldsymbol{\theta}} \,|\, \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$ constitutes a statistical model which merges the advantages of graphical Markov modelling with the distributional flexibility of the pair-copula approach. We will refer to such a model as a *pair-copula Bayesian network* (PCBN). We want to mention that PCBNs were first introduced in Kurowicka and Cooke (2005). The analyses therein were, however, restricted to pair-copula families with the property that zero rank correlation implies independence.

Simulation

Write $V = \{v_1, \ldots, v_d\}$ according to the well-ordering of $\mathcal{D}$ and set $V_{-i} := \{v_1, \ldots, v_{i-1}\}$ for all $i \in \{1, \ldots, d\}$. A sample $\boldsymbol{u} = (u_{v_1}, \ldots, u_{v_d}) \in [0,1]^d$ from a fully specified PCBN with uniform $[0,1]$ univariate margins is obtained by simulating d independent uniform $[0,1]$ variables

$x_1, \dots, x_d$ and applying the quantile transformations

$$\begin{aligned} u_{v_1} &:= x_1, \\ u_{v_2} &:= F^{-1}_{v_2|v_1}(x_2 \,|\, u_{v_1}; \boldsymbol{\theta}), \\ u_{v_3} &:= F^{-1}_{v_3|V_{-3}}(x_3 \,|\, \boldsymbol{u}_{V_{-3}}; \boldsymbol{\theta}), \\ &\vdots \\ u_{v_d} &:= F^{-1}_{v_d|V_{-d}}(x_d \,|\, \boldsymbol{u}_{V_{-d}}; \boldsymbol{\theta}). \end{aligned}$$

The order in which the components of $\boldsymbol{u}$ are generated is given by the well-ordering of $\mathcal{D}$. Solving transformation equation i for x_i, we have by the local $\mathcal{D}$-Markov property in Equation (2.1)

$$x_i = F_{v_i|V_{-i}}(u_{v_i} \,|\, \boldsymbol{u}_{V_{-i}}; \boldsymbol{\theta}) = F_{v_i|\mathrm{pa}(v_i)}(u_{v_i} \,|\, \boldsymbol{u}_{\mathrm{pa}(v_i)}; \boldsymbol{\theta}), \quad i \in \{1, \dots, d\}. \tag{4.9}$$

Now assume that $\mathrm{pa}(v_i) \neq \emptyset$, and let w denote the largest vertex in $\mathrm{pa}(v_i)$ by the parent ordering $<_{v_i}$. Then Equations (4.9) and (3.2) yield

$$x_i = h_{\underline{v_i}, w|\mathrm{pa}(v_i;w)}\big(F_{v_i|\mathrm{pa}(v_i;w)}(u_{v_i} \,|\, \boldsymbol{u}_{\mathrm{pa}(v_i;w)}; \boldsymbol{\theta}), F_{w|\mathrm{pa}(v_i;w)}(u_w \,|\, \boldsymbol{u}_{\mathrm{pa}(v_i;w)}; \boldsymbol{\theta}); \boldsymbol{\theta}\big). \tag{4.10}$$

Since u_{v_i} is only contained in the first h-function argument $F_{v_i|\mathrm{pa}(v_i;w)}(u_{v_i} \,|\, \boldsymbol{u}_{\mathrm{pa}(v_i;w)}; \boldsymbol{\theta})$ on the right hand side of Equation (4.10), we obtain by induction that the only inverse functions needed in the computation of u_{v_i} are the inverse h-functions $h^{-1}_{\underline{v_i}, w^*|\mathrm{pa}(v_i, w^*)}$, $w^* \in \mathrm{pa}(v_i)$.

ML estimation

ML estimation for PCBNs was first considered in Bauer et al. (2012). Let $\boldsymbol{u} = \big(\boldsymbol{u}^1, \dots, \boldsymbol{u}^n\big)$, $n \in \mathbb{N}$, be a realisation of a sample of i.i.d. observations $\boldsymbol{U}^1, \dots, \boldsymbol{U}^n$ from a random variable $\boldsymbol{U}$ on $[0,1]^d$ with copula family $\{C_{\mathcal{D},\mathcal{O},\boldsymbol{\theta}} \,|\, \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$ and uniform univariate margins. The restriction to uniform univariate margins is made along the same lines as in Section 3.2 for vine copula models. Equation (4.1) yields the log-likelihood function

$$l(\boldsymbol{\theta}; \boldsymbol{u}) = \sum_{k=1}^{n} \sum_{v \in V} \sum_{w \in \mathrm{pa}(v)} \log c_{v,w|\mathrm{pa}(v;w)}\Big(F_{v|\mathrm{pa}(v;w)}\big(u_v^k \,|\, \boldsymbol{u}^k_{\mathrm{pa}(v;w)}; \boldsymbol{\theta}\big), F_{w|\mathrm{pa}(v;w)}\big(u_w^k \,|\, \boldsymbol{u}^k_{\mathrm{pa}(v;w)}; \boldsymbol{\theta}\big); \boldsymbol{\theta}\Big). \tag{4.11}$$

ML estimation of the parameters in Equation (4.11) can be performed using a stepwise approach similar to the one discussed in Section 3.2 for vine copula models. The only difference to vine copula models is that we iterate over the vertices of $\mathcal{D}$ and their respective parents instead of over the trees of an R-vine. Hence again, in a first step, *sequential ML estimates* are computed and in a second step, using the sequential ML estimates as starting values, *joint ML estimates* $\widehat{\boldsymbol{\theta}}_{v,w|\mathrm{pa}(v;w)}$, $v \in V$, $w \in \mathrm{pa}(v)$, are inferred.

Model selection

Model selection for PCBNs involves estimation of the DAG $\mathcal{D}$, selection of the set $\mathcal{O}$ of parent orderings, and selection of the pair-copula families for $C_{v,w|\mathrm{pa}(v;w)}$, $v \in V$, $w \in \mathrm{pa}(v)$. Estimation of $\mathcal{D}$ will be the subject of Section 5. Given $\mathcal{D}$ and $\mathcal{O}$, the selection of pair-copula families can be performed in a similar way as in Section 3.2 for vine copula models, again with the difference that the iteration is vertex-by-vertex and parent-by-parent instead of tree-by-tree.

For the selection of $\mathcal{O}$ we propose a greedy-type procedure inspired by the structure selection algorithm for vine copula models outlined in Section 3.2. Clearly, an ordering of the parents of a vertex $v \in V$ is only required if $\mathrm{pa}(v) \neq \emptyset$. We assume $\mathcal{D}$ is well-ordered. Let $v \in V$ and assume that $k := |\mathrm{pa}(v)| \geq 1$. Moreover, let $i \in \{1, \dots, k\}$ and assume that we have already selected the $i-1$ smallest parents of v, denoted by $w_1 <_v \cdots <_v w_{i-1}$. This implies that we have already selected pair-copula families for $C_{v^*,w|\mathrm{pa}(v^*;w)}$, v^* smaller than v by the well-ordering of $\mathcal{D}$, $w \in \mathrm{pa}(v^*)$, and $C_{v,w_j|\mathrm{pa}(v;w_j)}$, $j < i$. Also, this implies that we have inferred corresponding ML parameter estimates, which we summarise in the vector $\widehat{\boldsymbol{\theta}}$. Let $W_{-i} := \{w_1, \dots, w_{i-1}\}$. The selection of w_i is performed in three steps. First, we compute the pseudo-observations $F_{v|W_{-i}}\big(u_v^k \,\big|\, \boldsymbol{u}_{W_{-i}}^k; \widehat{\boldsymbol{\theta}}\big)$ and $F_{w|W_{-i}}\big(u_w^k \,\big|\, \boldsymbol{u}_{W_{-i}}^k; \widehat{\boldsymbol{\theta}}\big)$, $k \in \{1, \dots, n\}$, for all $w \in \mathrm{pa}(v) \setminus W_{-i}$. Note that for $i = 1$, nothing needs to be done since all univariate marginals are uniform on $[0, 1]$. Second, we assign a weight $\omega_{v,w}$ to every edge $w \to v$, $w \in \mathrm{pa}(v) \setminus W_{-i}$, based on the previously calculated pseudo-observations, and choose w_i such that $w_i \to v$ has optimal edge weight. Suitable weights are, for instance, the absolute values of estimates of Kendall's τ, or AIC or BIC values of selected pair-copula families with estimated parameters. Last, we select a pair-copula family for $C_{v,w_i|\mathrm{pa}(v;w_i)}$ and compute an ML estimate of the corresponding parameter(s). Again, this last step may have already been performed when computing the edge weights $\omega_{v,w}$.

5. Structure estimation in Bayesian networks using the PC algorithm

The first task of modelling the joint distribution of a given set of variables with a Bayesian network is to identify the DAG $\mathcal{D} = (V, E)$ specifying the Markov structure of the variables. A convenient approach to defining $\mathcal{D}$ is the use of *expert knowledge*. However, the scope of this approach is rather limited since expert knowledge is often incomplete or unavailable. Data-driven *structure estimation algorithms* provide a computer-based alternative to elicited expert knowledge. Robinson (1973) has shown that the number n_d of DAGs on $d := |V|$ labelled vertices is given by the recurrence equation

$$n_0 = 1, \qquad n_d = \sum_{k=1}^{d} (-1)^{k-1} \binom{d}{k} 2^{k(d-k)} \, n_{d-k}.$$

Since n_d grows super-exponentially in d, a systematic trial of all possible DAGs on V is infeasible, and thus efficient searching algorithms are required. A considerable number of structure estimation algorithms has been proposed over the last two decades, see Neapolitan (2003, Chapters 8 – 11) and Koller and Friedman (2009, Chapter 18) for an overview. The majority of these algorithms follow one of the two estimation approaches predominant in the literature: the *constraint-based* and the *score-and-search-based* approach. In the constraint-based approach, $\mathcal{D}$ is inferred from a series of conditional independence tests. In the score-and-search-based approach, $\mathcal{D}$ is found by optimising a given scoring function—like AIC or BIC—over a suitable search space, for instance the space of all DAGs or the space of all Markov-equivalence classes. Besides, there exist hybrid algorithms which combine both approaches. Unfortunately, available implementations of aforementioned algorithms are mainly confined to discrete or Gaussian models and are hence not suited for our non-Gaussian continuous Bayesian networks.

5.1. The PC algorithm

We will provide a structure estimation algorithm that is particularly suited to finding the DAG $\mathcal{D} = (V, E)$ underlying a non-Gaussian continuous Bayesian network. Our algorithm is a version of one of the most popular constraint-based estimation algorithms, the *PC algorithm* (named after its inventors *P*eter Spirtes and *C*lark Glymour), see Spirtes and Glymour (1991) and Spirtes et al. (2000, Section 5.4.2). To fix notation and for the reader's convenience, we will now recall the PC algorithm. Let P be an absolutely continuous $\mathcal{D}$-Markovian probability measure on $[0,1]^d$ with uniform univariate margins. The restriction to uniform univariate margins is made along the same lines as in Section 4.2. Moreover, let $\boldsymbol{u} = \big(\boldsymbol{u}^1, \ldots, \boldsymbol{u}^n\big)$, $n \in \mathbb{N}$, be a realisation of a sample of i.i.d. observations $\boldsymbol{U}^1, \ldots, \boldsymbol{U}^n$ from a random variable $\boldsymbol{U}$ distributed as P. The PC algorithm for estimating $\mathcal{D}$ from $\boldsymbol{u}$ involves three major steps in which the complete UG $\mathcal{G}$ on V is gradually transformed into a CG $\mathcal{G}^*$ on V, which is supposed to be the essential graph $\mathcal{D}^e$ corresponding to the Markov-equivalence class $[\mathcal{D}]$ of $\mathcal{D}$. The resulting CG $\mathcal{G}^*$ can then be extended to a DAG as outlined in Section 2.

In the first step of the PC algorithm, a series of tests for conditional independence is performed on $\boldsymbol{u}$. More precisely, for all distinct vertices $i, j \in V$ and chosen vertex sets $K \subseteq V \setminus \{i, j\}$, the null hypothesis $\mathrm{H}_0 \colon U_i \perp\!\!\!\perp U_j \mid \boldsymbol{U}_K$ is tested against the general alternative $\mathrm{H}_1 \colon U_i \not\perp\!\!\!\perp U_j \mid \boldsymbol{U}_K$ of conditional dependence. Given a suitable independence test of choice, we denote the test decision at significance level $\alpha \in (0,1)$ by $T_\alpha(\boldsymbol{u}_i, \boldsymbol{u}_j; \boldsymbol{u}_K) \in \{\mathrm{H}_0, \mathrm{H}_1\}$. We will later introduce a novel class of conditional independence tests that is particularly tailored to the algorithm and applicable to non-Gaussian continuous data. If $T_\alpha(\boldsymbol{u}_i, \boldsymbol{u}_j; \boldsymbol{u}_K) = \mathrm{H}_0$, the edge $i - j$ is removed from $\mathcal{G}$ and the conditioning set K is stored in two variables S_{ij} and S_{ji} for later use. As a result of the first step, $\mathcal{G}$ is turned into the skeleton of $\mathcal{G}^*$. Step one is given in Algorithm 3.

Algorithm 3 PC algorithm: finding the skeleton.

Input Data set $\boldsymbol{u}$; significance level $\alpha \in (0,1)$; conditional independence test with test decision $T_\alpha(\boldsymbol{u}_i, \boldsymbol{u}_j; \boldsymbol{u}_K)$ for the null hypothesis $\mathrm{H}_0\colon U_i \perp\!\!\!\perp U_j \mid \boldsymbol{U}_K$, $i \neq j \in V$, $K \subseteq V \setminus \{i,j\}$.

Output Skeleton $\mathcal{G} = (V, E_\mathcal{G})$; separation sets S_{ij}, $i \neq j \in V$, $(i,j) \notin E_\mathcal{G}, (j,i) \notin E_\mathcal{G}$.

```
G ← complete UG on V;
k ← 0;
repeat
    for i ∈ V and j ∈ ad(i) do % i and j are adjacent in G
        if T_α(u_i, u_j; u_K) = H_0 for any K ⊆ ad(i) \ {j} with |K| = k then
            delete i − j from G;
            S_ij ← K;
            S_ji ← K;
        end if
    end for
    k ← k + 1.
until |ad(i)| ≤ k for all i ∈ V.
```

In the second step, $\mathcal{G}$ is transformed into a CG by introducing a v-structure $i \to k \leftarrow j$ whenever i and j are non-adjacent, $k \in \mathrm{ad}(i) \cap \mathrm{ad}(j)$, and $k \notin S_{ij}$. In the last step, $\mathcal{G}$ is transformed into $\mathcal{G}^*$ by directing further edges of $\mathcal{G}$ to prevent new v-structures and directed cycles, until no more edges need direction. Steps two and three are given in Algorithm 4, where the third step was taken from Pearl (2009, Section 2.5). If P is faithful to $\mathcal{D}$ and if all statistical test decisions made in Algorithm 3 are correct, then Algorithm 4 will return the correct graph $\mathcal{D}^e$, see Meek (1995). Due to the finite sample size or the existence of hidden variables, the application of Algorithm 3 to empirical data may sometimes, however, lead to conflicting information about edge directions. That is, it may be possible in a given situation that Algorithm 4, while introducing v-structures, first orients an undirected edge $i - j$ into $i \to j$, and later tries to introduce $i \leftarrow j$. In such a situation, we keep $i \to j$ and skip the new v-structure including $i \leftarrow j$. We can test whether the resulting CG can still be extended to a DAG without introducing new v-structures or directed cycles using the algorithm by Dor and Tarsi (1992). The PC algorithm can also be adapted to incorporate existing expert knowledge, see Meek (1995) and Moole and Valtorta (2004). We will henceforth assume that P is faithful to $\mathcal{D}$ and that there are no hidden variables.

Testing conditional independence using partial correlations

The centrepiece of the PC algorithm—as of any constraint-based estimation algorithm—is the test for conditional independence. In a Gaussian framework, the test of choice is usually a test for zero partial correlation $\rho_{ij\cdot K}$, see, for instance, Anderson (2003, Section 4.3). The null hypothesis then translates into $\mathrm{H}_0\colon \rho_{ij\cdot K}(X_i, X_j; X_K) = 0$, where $X_k := \Phi^{-1}(U_k)$ for all $k \in V$, and Φ denotes the univariate standard normal cdf. Here, the quantile function Φ^{-1} is applied to U in order to transform the uniform univariate copula margins to standard normal margins.

Algorithm 4 PC algorithm: introducing edge directions

Input Skeleton $\mathcal{G} = (V, E_{\mathcal{G}})$; separation sets S_{ij}, $i \neq j \in V$, $(i,j) \notin E_{\mathcal{G}}, (j,i) \notin E_{\mathcal{G}}$.
Output Chain graph $\mathcal{G}$.

```
% Introduce v-structures:
for i ∈ V and j ∉ ad(i) and k ∈ ad(i) ∩ ad(j) do
    if k ∉ S_ij then
        replace i − k − j by i → k ← j in G;
    end if
end for
% Orient as many undirected edges as possible by repeated application of the following rules:
repeat
    R1 orient j − k into j → k whenever G contains i → j and k ∉ ad(i);
    R2 orient i − j into i → j whenever G contains i → k → j;
    R3 orient i − j into i → j whenever G contains i − k → j and i − l → j, and l ∉ ad(k);
until no more edges can be directed;
```

The conditional independence test is based on the asymptotic normality

$$\sqrt{d - |K| - 3}\, \widehat{z}_n \xrightarrow[n\to\infty]{\mathcal{L}} \mathrm{N}(0,1), \qquad \widehat{z}_n := \frac{1}{2} \log\left(\frac{1 + \widehat{\rho}_{ij\cdot K}(\boldsymbol{X}_i^n, \boldsymbol{X}_j^n; \boldsymbol{X}_K^n)}{1 - \widehat{\rho}_{ij\cdot K}(\boldsymbol{X}_i^n, \boldsymbol{X}_j^n; \boldsymbol{X}_K^n)}\right),$$

of the Fisher's z-transformed partial-correlation estimator $\widehat{\rho}_{ij\cdot K}$ under H_0, see again Anderson (2003, Section 4.3). Here, $\xrightarrow{\mathcal{L}}$ denotes convergence in distribution, $\mathrm{N}(0,1)$ is the univariate standard normal distribution, and $\boldsymbol{X}_k^n := \big(\Phi^{-1}(U_k^1), \ldots, \Phi^{-1}(U_k^n)\big)$ for all $k \in V$. Kalisch and Bühlmann (2007) have proven uniform convergence of the PC algorithm under joint normality and a mild sparsity assumption for the underlying DAG, cf. also Harris and Drton (2012). An implementation of the PC algorithm with above partial correlation test is available in the R package `pcalg` (Kalisch et al., 2012). The `pcalg` package also provides an interface for self-implemented conditional independence tests.

5.2. Testing conditional independence using vine copulas and the Rosenblatt transform

Above test for zero partial correlation was derived under the assumption of joint normality. We will now introduce a copula-based alternative test for conditional independence that is also applicable to non-Gaussian continuous data. Assume $K \neq \emptyset$. Otherwise, the problem reduces to testing ordinary (unconditional) stochastic independence. Let $F_{i,j|K}(\,\cdot\,,\cdot \mid \boldsymbol{v}_K)$ denote the conditional cdf of U_i and U_j given $\boldsymbol{U}_K = \boldsymbol{v}_K$, and let $C_{i,j|K}(\,\cdot\,,\cdot \mid \boldsymbol{v}_K)$ be the corresponding conditional copula. Moreover, let $C_{\perp\!\!\!\perp}$ denote the independence copula on $[0,1]^2$. The conditional independence $U_i \perp\!\!\!\perp U_j \mid \boldsymbol{U}_K$ holds if and only if

$$F_{i,j|K}(v_i, v_j \mid \boldsymbol{v}_K) = C_{i,j|K}\big(F_{i|K}(v_i \mid \boldsymbol{v}_K), F_{j|K}(v_j \mid \boldsymbol{v}_K) \,\big|\, \boldsymbol{v}_K\big) = F_{i|K}(v_i \mid \boldsymbol{v}_K)\, F_{j|K}(v_j \mid \boldsymbol{v}_K)$$

for all $v_i, v_j \in [0,1]$ and P_K-almost all $\boldsymbol{v}_K \in [0,1]^{|K|}$, where $\boldsymbol{U}_K \sim P_K$. Hence, the null hypothesis of the conditional independence test can be stated as $H_0\colon C_{i,j|K}(\,\cdot\,,\cdot\,|\,\boldsymbol{v}_K) = C_{\perp\!\!\!\perp}(\,\cdot\,,\cdot\,)$ for P_K-almost all $\boldsymbol{v}_K \in [0,1]^{|K|}$. Using the simplifying assumption that $C_{i,j|K}(\,\cdot\,,\cdot\,|\,\boldsymbol{v}_K)$ depends on $\boldsymbol{v}_K$ only through $F_{i|K}(\,\cdot\,,\cdot\,|\,\boldsymbol{v}_K)$ and $F_{j|K}(\,\cdot\,,\cdot\,|\,\boldsymbol{v}_K)$ discussed in Section 3.2, we drop $\boldsymbol{v}_K$ from $C_{i,j|K}(\,\cdot\,,\cdot\,|\,\boldsymbol{v}_K)$ and approximate H_0 by the more accessible null hypothesis $H_0^*\colon C_{i,j|K}(\,\cdot\,,\cdot\,) = C_{\perp\!\!\!\perp}(\,\cdot\,,\cdot\,)$. The new null hypothesis H_0^* can be tested using any test for ordinary (unconditional) stochastic independence of two continuous random variables applied to the transformed observations $W_{i|K}^1, \ldots, W_{i|K}^n$ and $W_{j|K}^1, \ldots, W_{j|K}^n$, where

$$W_{i|K}^k \coloneqq F_{i|K}\big(U_i^k \,\big|\, \boldsymbol{U}_K^k\big) \quad \text{and} \quad W_{j|K}^k \coloneqq F_{j|K}\big(U_j^k \,\big|\, \boldsymbol{U}_K^k\big) \tag{5.1}$$

for all $k \in \{1, \ldots, n\}$. Song (2009) called Equation (5.1) the *Rosenblatt transform* after Rosenblatt (1952), while Bergsma (2011) called it the *partial copula transform*. Given a realisation $\boldsymbol{u}$ of $\big(\boldsymbol{U}^1, \ldots, \boldsymbol{U}^n\big)$, the difficulty of this approach lies in the computation of the transformed realisations $\boldsymbol{w}_{i|K}$ and $\boldsymbol{w}_{j|K}$, where $w_{i|K}^k \coloneqq F_{i|K}\big(u_i^k \,\big|\, \boldsymbol{u}_K^k\big)$ and $w_{j|K}^k \coloneqq F_{j|K}\big(u_j^k \,\big|\, \boldsymbol{u}_K^k\big)$ for all $k \in \{1, \ldots, n\}$. Note that the conditional cdfs $F_{i|K}(\,\cdot\,|\,\boldsymbol{v}_K)$ and $F_{j|K}(\,\cdot\,|\,\boldsymbol{v}_K)$ are typically unknown and need to be estimated in the course of the testing procedure. Bergsma (2011) suggested the use of non-parametric kernel estimators for this task. By contrast, we propose a parametric estimation method that is based on vine copula models.

Estimating conditional cdfs using vine copula models

Taking another look at vine copula models as described in Section 3.2, we observe that transformed realisations like $\boldsymbol{w}_{i|K}$ and $\boldsymbol{w}_{j|K}$ naturally emerge in the log-likelihood function. In fact, given any distinct $i, j \in V$ and $K \subseteq V \setminus \{i,j\}$, it is always possible to construct a regular vine $\mathcal{V} = (T_1, \ldots, T_p)$, $p \coloneqq 1 + |K|$, in which tree T_1 has vertex set $V_1 = \{i\} \cup \{j\} \cup K$ and tree T_p is of the form $i, l|K_{-l} \overset{i,j|K}{\text{---}} j, m|K_{-m}$ for some $l, m \in K$. The corresponding log-likelihood function $l\big(\boldsymbol{\theta}; \boldsymbol{u}_{\{i\}\cup\{j\}\cup K}\big)$, $\boldsymbol{\theta} \in \Theta$, contains the pair-copula pdf $c_{i,j|K}$ with arguments $F_{i|K}\big(u_i^k \,\big|\, \boldsymbol{u}_K^k; \boldsymbol{\theta}\big)$ and $F_{j|K}\big(u_j^k \,\big|\, \boldsymbol{u}_K^k; \boldsymbol{\theta}\big)$ for all $k \in \{1, \ldots, n\}$. Thus, by computing an ML estimate $\widehat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ and subsequently evaluating l at $\widehat{\boldsymbol{\theta}}$, we obtain estimates $\widehat{w}_{i|K}^k \coloneqq F_{i|K}\big(u_i^k \,\big|\, \boldsymbol{u}_K^k; \widehat{\boldsymbol{\theta}}\big)$ and $\widehat{w}_{j|K}^k \coloneqq F_{j|K}\big(u_j^k \,\big|\, \boldsymbol{u}_K^k; \widehat{\boldsymbol{\theta}}\big)$ of $w_{i|K}^k$ and $w_{j|K}^k$, respectively, as a welcome side effect.

We call a vertex v in Tree T_q, $q \in \{1, \ldots, p-1\}$, an *inner vertex* if $|\mathrm{ad}(v)| \geq 2$. In order to construct such a vine $\mathcal{V}$, we have to follow one simple rule:

R Neither i nor j may be part of an inner vertex in the trees $T_1, \ldots, T_{p-1}$ of $\mathcal{V}$.

Following **R**, it is even possible to restrict the class of R-vines to C- or D-vines. The only inner vertices of a C-vine are the root vertices of the trees $T_1, \ldots, T_{p-1}$. Thus, in a C-vine obeying **R**, i and j do not appear in the root vertices of the respective trees. Similarly, in a D-vine obeying

R, i and j only appear in the boundary vertices of trees $T_1, \ldots, T_{p-1}$. Figures 2 and 5 give an example of a C-, a D-, and an R-vine, respectively, having the same edge label in tree T_p.

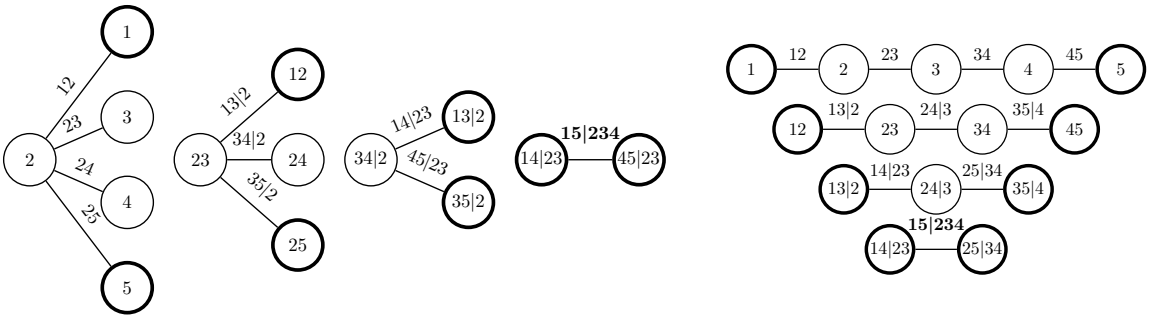

Figure 5: A C- (left) and a D-vine (right) on five vertices having the same edge label 15|234 in tree T_4. A corresponding R-vine is given in Figure 2. The three vines were constructed according to rule **R** with $i = 1$ and $j = 5$. Boundaries of nodes including either 1 or 5 appear in bold.

The tree structure of $\mathcal{V}$ can be estimated from $\boldsymbol{u}_{\{i\}\cup\{j\}\cup K}$ by adapting the greedy search strategies described in Section 3.2 to the new constraint **R**. An optimal C-vine obeying **R** is found by restricting the sets of possible root vertices for trees $T_1, \ldots, T_{p-1}$ to vertices containing neither i nor j, respectively. In order to find an optimal D-vine obeying **R**, the unconstrained TSP usually solved has to be replaced by a constrained TSP with fixed source vertex i and destination vertex j. Finally, an optimal R-vine obeying **R** is found by first estimating a smaller R-vine $\mathcal{V}_K$ with first tree vertices K. Having found $\mathcal{V}_K$, vertex i is then connected to a vertex $l \in K$ in tree T_1 such that the new edge $i - l$ has optimal edge weight amongst all possible edges $i - m$ for $m \in K$. The same is done for vertex j. Note that this way, j cannot be connected to i. The newly formed structure is then sequentially transformed into $\mathcal{V}$ by analogously extending the remaining trees $T_2, \ldots, T_p$, such that the proximity condition and **R** are always satisfied and the corresponding edge weights are optimised. Copula selection and ML estimation in the resulting vine copula model is then performed as usual, see Section 3.2.

Vine-copula-based conditional independence tests

Summing up, we test the conditional independence $U_i \perp\!\!\!\perp U_j \mid \boldsymbol{U}_K$ in three steps. In the first step, we construct a vine $\mathcal{V}$ on the vertices $\{i\}\cup\{j\}\cup K$ by applying a modified version of one of the structure estimation algorithms described in Section 3.2 to $\boldsymbol{u}_{\{i\}\cup\{j\}\cup K}$. In the second step, we select corresponding pair-copula families, perform ML estimation in the resulting model, and evaluate the log-likelihood function l at the estimated parameter vector $\widehat{\boldsymbol{\theta}}$ to obtain transformed realisations $\widehat{\boldsymbol{w}}_{i|K} := \left(\widehat{w}^k_{i|K}\right)_{1\le k\le n}$ and $\widehat{\boldsymbol{w}}_{j|K} := \left(\widehat{w}^k_{j|K}\right)_{1\le k\le n}$, respectively. In the last step, we apply a test for ordinary stochastic independence of two continuous random variables to $\widehat{\boldsymbol{w}}_{i|K}$

and $\widehat{\boldsymbol{w}}_{j|K}$. Note that in the first iteration step of Algorithm 3, only unconditional independences, that is $K = \emptyset$, are tested, and thus the independence test of choice is directly applied to $\boldsymbol{u}$.

We will examine the performance of our novel testing procedure in a simulation study in Section 6 using three different tests for ordinary stochastic independence. Recycling notation, consider the null hypothesis $H_0 \colon U_i \perp\!\!\!\perp U_j$ vs. $H_1 \colon U_i \not\perp\!\!\!\perp U_j$. The first test used is a test for zero Kendall's τ with null hypothesis $H_0^* \colon \tau(U_i, U_j) = 0$ vs. $H_1^* \colon \tau(U_i, U_j) \neq 0$. Under H_0, the Kendall's τ estimator $\widehat{\tau}_n$ exhibits the asymptotic normality

$$\sqrt{\frac{9n(n-1)}{2(2n+5)}}\, \widehat{\tau}_n(\boldsymbol{U}_i, \boldsymbol{U}_j) \xrightarrow[n\to\infty]{\mathcal{L}} \mathrm{N}(0,1),$$

where $\boldsymbol{U}_i \coloneqq \left(U_i^1, \ldots, U_i^n\right)$ and $\boldsymbol{U}_j \coloneqq \left(U_j^1, \ldots, U_j^n\right)$, see Hollander and Wolfe (1999, Section 8.1). In general, $\tau(U_i, U_j) = 0$ does not imply $U_i \perp\!\!\!\perp U_j$. However, for many popular copula families like the Clayton, the Gaussian, and the Gumbel copula families, H_0 and H_0^* are equivalent. The family of Student's t copulas serves as a counterexample. We then consider H_0^* an approximation for H_0. The other two independence tests used in Section 6 are of Cramér-von Mises type. More precisely, independence test number two is the test for zero Hoeffding's D proposed by Hoeffding (1948). P-values of the sample test statistic $\widehat{D}_n$ are computed using the asymptotically equivalent sample test statistic $\widehat{B}_n$ by Blum et al. (1961), see also Hollander and Wolfe (1999, Section 8.6). Independence test number three is the test by Genest and Rémillard (2004) based on the empirical copula process.

6. Simulation study

We conducted an extensive simulation study to examine the small sample performance of the PC algorithm in finding the true Markov structure underlying a PCBN. To this end, we drew samples from various PCBNs based on the conditional independence properties represented by the DAG $\mathcal{D} = (V, E)$ in Figure 1. These PCBNs emerged from various choices of pair-copula families for C_{12}, C_{13}, C_{24}, and $C_{34|2}$, cf. Section 4. More precisely, we chose from the Clayton, Gumbel, Gaussian, and Student's t pair-copula families. These copula families exhibit considerable differences in their dependence structures and tail behaviours, see the simulation study in Bauer et al. (2012) for an overview. We considered four PCBNs with all four pair copulas C_{12}, C_{13}, C_{24}, and $C_{34|2}$ coming from the same copula family, respectively. Additionally, we considered 24 PCBNs with each pair copula C_{12}, C_{13}, C_{24}, and $C_{34|2}$ coming from a different copula family. Our choices of pair-copula families are given in Table 2. For each choice of pair-copula families we then considered 16 different parameter configurations arising from a selection of two different parameter values for each pair copula. The parameter values for each pair copula were chosen to correspond to values of Kendall's τ of 0.25 and 0.75, that is one low and one high

rank-correlation specification. These Kendall's τ configurations are summarised in Table 3.

Copula	1	2	3	4	5	6	7	8	9	10	11	12	13	14
C_{12}	C	G	N	t	C	C	C	C	C	C	G	G	G	G
C_{13}	C	G	N	t	G	G	N	N	t	t	C	C	N	N
C_{24}	C	G	N	t	N	t	G	t	G	N	N	t	C	t
$C_{34\|2}$	C	G	N	t	t	N	t	G	N	G	t	N	t	C
Copula	15	16	17	18	19	20	21	22	23	24	25	26	27	28
C_{12}	G	G	N	N	N	N	N	N	t	t	t	t	t	t
C_{13}	t	t	C	C	G	G	t	t	C	C	G	G	N	N
C_{24}	C	N	G	t	C	t	C	G	G	N	C	N	C	G
$C_{34\|2}$	N	C	t	G	t	C	G	C	N	G	N	C	G	C

Table 2: Selected pair-copula families for C_{12}, C_{13}, C_{24}, $C_{34|2}$. Copulas were chosen from the Clayton (C), Gumbel (G), Gaussian (N), and Student's t (t) pair-copula families. See Tables 3 and 4 for further details on the pair-copula families used.

Copula	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
C_{12}	0.25	0.75	0.25	0.25	0.25	0.75	0.75	0.75	0.25	0.25	0.25	0.75	0.75	0.75	0.25	0.75
C_{13}	0.25	0.25	0.75	0.25	0.25	0.75	0.25	0.25	0.75	0.75	0.25	0.75	0.75	0.25	0.75	0.75
C_{24}	0.25	0.25	0.25	0.75	0.25	0.25	0.75	0.25	0.75	0.25	0.75	0.75	0.25	0.75	0.75	0.75
$C_{34\|2}$	0.25	0.25	0.25	0.25	0.75	0.25	0.25	0.75	0.25	0.75	0.75	0.25	0.75	0.75	0.75	0.75

Table 3: Selected values of Kendall's τ for each choice of pair-copula families for C_{12}, C_{13}, C_{24}, $C_{34|2}$. See Tables 2 and 4 for further details on the pair-copula families used.

Our selection of copula parameters is based on the bijective relationship between the parameters of the Clayton, Gumbel, and Gaussian pair-copula families and the corresponding Kendall's τ. For the Student's t copula, such a bijective relationship exists only between the correlation parameter and Kendall's τ, which is why we set the degrees-of-freedom parameter of each Student's t copula to $\nu = 5$ in order to allow for heavy-tailed dependence. Table 4 summarises the parameters θ, the corresponding Kendall's correlation coefficients $\tau(\theta)$, and the respective tail-dependence coefficients $\lambda_{\mathrm{L}}(\theta) = \lim_{u\to 0} \frac{C_\theta(u,u)}{u}$ and $\lambda_{\mathrm{U}}(\theta) = \lim_{u\to 1} \frac{1-2u+C_\theta(u,u)}{1-u}$ for each pair copula C_θ, $\theta \in \Theta$, used in the simulation study.

Summing up, we have 28 different PCBNs with 16 different parameter configurations each, that is 448 simulation scenarios. In each of the 448 simulation scenarios we performed $N = 100$ simulation runs, and in each simulation run we generated $n = 1{,}000$ i.i.d. observations. The sampling procedure used was described in Section 4.2.

For each of the 44,800 runs we applied the PC algorithm with the ten different conditional independence tests described in Section 5. Those were the widely used test for zero partial

Copula	Clayton		Gumbel		Gauss		Student	
Parameter(s)	0.67	6.00	1.33	4.00	0.38	0.92	0.38, 5	0.92, 5
Kendall's τ	0.25	0.75	0.25	0.75	0.25	0.75	0.25	0.75
Lower TDC λ_{L}	0.35	0.89	0.00	0.00	0.00	0.00	0.15	0.15
Upper TDC λ_{U}	0.00	0.00	0.32	0.81	0.00	0.00	0.64	0.64

Table 4: Parameters, Kendall's correlation coefficients, and tail-dependence coefficients (TDCs) of the pair copulas used in the simulation study.

correlation (COR) and our novel vine-copula-based tests using either only C-vines (C), or only D-vines (D), or more generally R-vines (R), respectively, together with one of the Kendall's τ (K), Hoeffding's D (H), or Genest and Rémillard (GR) tests for ordinary stochastic independence. Since zero partial correlation is generally a weaker property than conditional independence, we consider COR only an approximate conditional independence test serving as a benchmark. In a Gaussian framework, however, zero partial correlation is equivalent to conditional independence. This equivalence holds in particular in the scenarios featuring only Gaussian pair copulas, in which case the respective joint copula families are also Gaussian. The corresponding correlation matrices were derived in Bauer et al. (2012). Each test was performed at the 5% significance level.

Results

Let $\mathcal{G}_{f,p,r,t}$ denote the CG obtained from applying the PC algorithm with conditional-independence test $t \in \{\text{COR}, \text{C-GR}, \text{C-H}, \text{C-K}, \text{D-GR}, \text{D-H}, \text{D-K}, \text{R-GR}, \text{R-H}, \text{R-K}\}$ to the data simulated in run $r \in \{1, \ldots, 100\}$ of pair-copula scenario $f \in \{1, \ldots, 28\}$ (see Table 2) and parameter configuration $p \in \{1, \ldots, 16\}$ (see Tables 3 and 4). We compared each CG $\mathcal{G}_{f,p,r,t}$ to the true essential graph $\mathcal{D}^e$ in Figure 1, and set $\pi_{f,p,r,t} := 1$ if $\mathcal{G}_{f,p,r,t}$ equalled $\mathcal{D}^e$ and $\pi_{f,p,r,t} := 0$ otherwise. For each pair-copula scenario f and each conditional independence test t, we then computed the relative frequency of recovering the correct structure over all parameter configurations p and all runs r, which we will denote by $\pi_{f,t} := \frac{1}{1600} \sum_{p=1}^{16} \sum_{r=1}^{100} \pi_{f,p,r,t}$. Moreover, we determined the *structural Hamming distance* (SHD) (Tsamardinos et al., 2006) $\delta_{f,p,r,t}$ between each CG $\mathcal{G}_{f,p,r,t}$ and $\mathcal{D}^e$. In short, $\delta_{f,p,r,t}$ counts the number of edges that need to be added to, removed from, directed in, or flipped in $\mathcal{G}_{f,p,r,t}$ in order to obtain $\mathcal{D}^e$. Hence, $\delta_{f,p,r,t}$ takes a value between zero and $\binom{|V|}{2} = 6$. We again took the average over all parameter configurations p and all runs r, yielding the mean SHD $\delta_{f,t} := \frac{1}{1600} \sum_{p=1}^{16} \sum_{r=1}^{100} \delta_{f,p,r,t}$ for each pair-copula scenario f and each conditional independence test t. The results are given in Figures 6 and 7, respectively.

Let us first consider Figure 6. The relative frequencies $\pi_{f,\text{COR}}$ range between 14% and 63%, whereas for the vine-copula-based tests, $\pi_{f,t}$ ranges between 40% and 64%. COR was outperformed by at least one vine-copula-based test in 18, and by all vine-copula-based tests in 15 out of the 28 copula scenarios. The lowest frequency of 14% was obtained when applying the PC

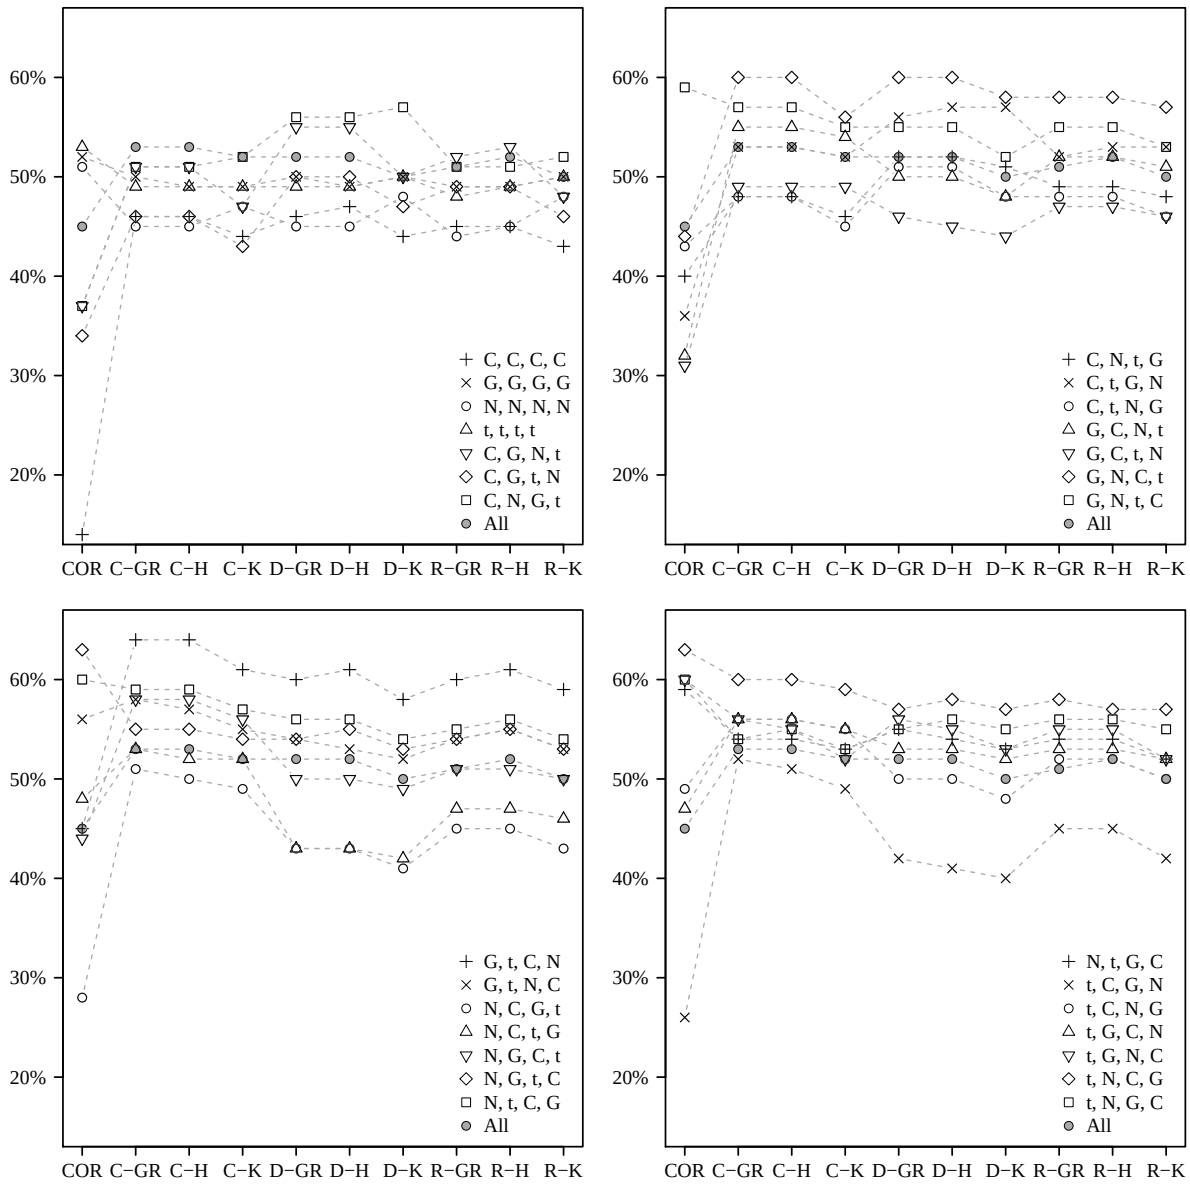

Figure 6: Percentage $\pi_{f,t}$ of runs in which the PC algorithm returned the correct Markov structure for each choice f of pair-copula families for C_{12}, C_{13}, C_{24}, $C_{34|2}$ (legends) and each conditional independence test t (horizontal axes) (1600 runs each). Copulas were chosen from the Clayton (C), Gumbel (G), Gaussian (N), and Student's t (t) pair-copula families. The percentage of correct recoveries out of all 28 copula scenarios is given in solid grey.

algorithm with COR to the data sets generated in copula scenario 1 (numbering as in Table 2), which features only Clayton, that is non-elliptical, copulas. By contrast, COR showed a solid performance in the elliptical-copulas-only scenarios 3 and 4, which is not surprising given that COR is based on the partial correlation. In 9 out of the 28 copula scenarios, $\pi_{f,\mathrm{COR}}$ is lower

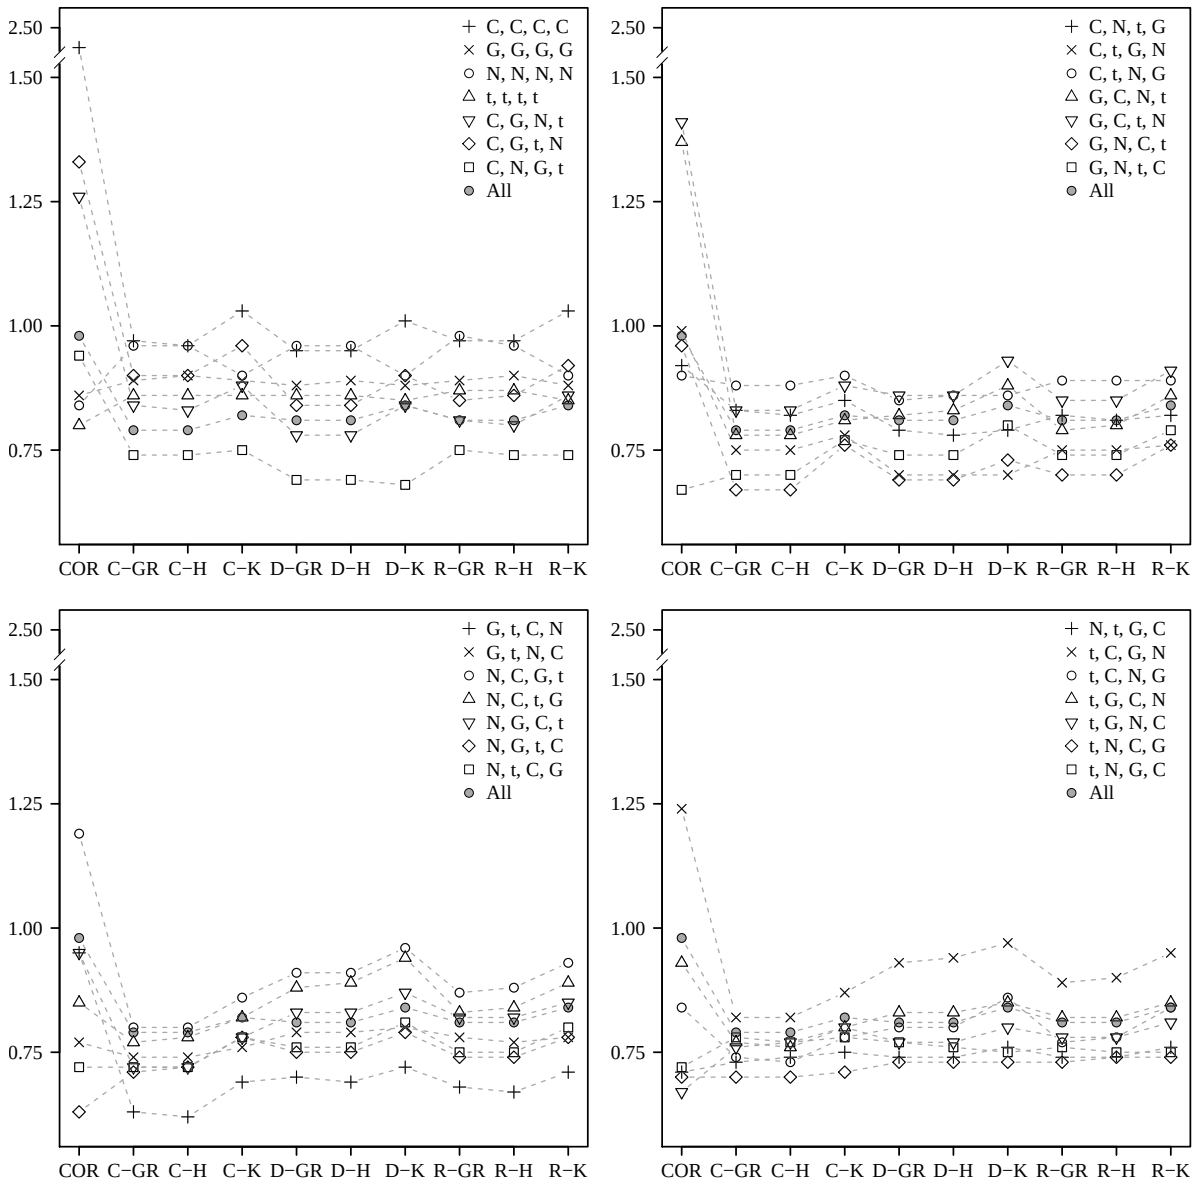

Figure 7: Average structural Hamming distance (SHD) $\delta_{f,t}$ between the true essential graph $\mathcal{D}^e$ and the chain graph $\mathcal{G}_{f,p,r,t}$ returned by the PC algorithm for each choice f of pair-copula families for C_{12}, C_{13}, C_{24}, $C_{34|2}$ (legends) and each conditional independence test t (horizontal axes) (1600 runs each). Copulas were chosen from the Clayton (C), Gumbel (G), Gaussian (N), and Student's t (t) pair-copula families. The average SHD over all 28 copula scenarios is given in solid grey.

than 40%, which is the minimum frequency obtained for the vine-copula-based tests. Also, in these 9 scenarios, the difference in relative frequencies between COR and the vine-copula-based tests ranges between 9 and 33 percentage points. The highest frequency of 64% was obtained in copula scenario 15 both for the PC algorithm with C-GR and C-H, respectively. Taking means

over all 28 copula scenarios, we obtain the overall relative frequencies $\pi_t := \frac{1}{28}\sum_{f=1}^{28}\pi_{f,t}$ for all tests t. These overall frequencies range between 50% and 53% for the vine-copula-based tests, while $\pi_{\text{COR}} = 45\%$. The best performances were again achieved by C-GR and C-H. However, we recommend using the R-vine-based conditional independence tests in higher dimensions since these offer more general tree structures than their C- and D-vine counterparts. Moreover, we observe that choosing H instead of GR as test for unconditional stochastic independence has only little effect on the performance of the vine-copula-based tests. By contrast, relative frequencies were, on average, slightly worse when using K instead of GR and H, respectively. Since zero Kendall's τ is generally also not equivalent to stochastic independence, we recommend using GR and H. Note that in a given copula scenario f and a given parameter scenario p, the relative frequencies $\pi_{f,p,t} := \frac{1}{100}\sum_{r=1}^{100}\pi_{f,p,r,t}$ can be a lot higher than the averages displayed in Figure 6. We observed frequencies $\pi_{f,p,t}$ of up to 98%. To sum up, using a vine-copula-based conditional independence test instead of COR leads to more reliable structure estimates, in particular when the data exhibit non-Gaussian, asymmetric dependence.

Considering only the correctly recovered Markov structures may be a too crude performance measure. Hence, the mean SHDs $\delta_{f,t}$ in Figure 7 illustrate how much the results of the PC algorithm differ from the true essential graph $\mathcal{D}^e$. For the vine-copula-based tests, $\delta_{f,t}$ ranges between 0.62 and 1.03. The respective overall means $\delta_t := \frac{1}{28}\sum_{f=1}^{28}\delta_{f,t}$ lie between 0.79 and 0.84. Thus, on average, the results of the PC algorithm differ by less than one edge from $\mathcal{D}^e$. That is, if the PC algorithm yields a CG that is not equivalent to $\mathcal{D}^e$, then, with a high probability, CG and $\mathcal{D}^e$ are not too different. The lowest values of δ_t were again obtained for C-GR and C-H. Similarly, $\delta_{f,\text{COR}}$ ranges between 0.63 and 2.44, and $\delta_{\text{COR}} = 0.98$, which again shows the superiority of the vine copula approach. The worst mean SHD of 2.44 was obtained in copula scenario 1. Overall, we can say that the PC algorithm with either of the 9 vine-copula-based conditional independence tests provides a suitable procedure for structure estimation in PCBNs.

We repeated the simulation study both for a significance level α of 1% and for a sample size n of 500. For $\alpha = 1\%$, we obtained results similar to the ones described above for $\alpha = 5\%$. The overall relative frequencies π_t were slightly lower, ranging from 44% to 47% for the vine-copula-based tests, while π_{COR} was 43%. Also, the overall mean SHDs δ_t ranged between 0.86 and 0.94 for the vine-copula-based tests, while π_{COR} was 0.99. The reduction in sample size to $n = 500$, on the other hand, lead to a slightly stronger decrease in the overall relative frequencies π_t, which then ranged between 39% and 41% for the vine-copula-based tests, while π_{COR} was 37%. Similarly, the overall mean SHDs δ_t ranged between 1.07 and 1.11 for the vine-copula-based tests, while π_{COR} was 1.17. Yet, both for $\alpha = 1\%$ and for $n = 500$, the CGs returned by the PC algorithm differed on average from $\mathcal{D}^e$ by only one edge. The performance of the PC algorithm can thus be deemed reliable and robust.

7. Application: Stock market indices

As a real-world application, we applied PCBNs to a financial data set comprising ten major international stock market indices. More precisely, we modelled the joint distribution of a portfolio of daily log-returns of the Australian All Ordinaries (AUS), the Canadian S&P/TSX Composite Index (CAN), the Swiss Market Index (CH), the German DAX (DEU), the French CAC 40 (FRA), the Hong Kong Hang Seng Index (HK), the Japanese Nikkei 225 (JPN), the Singapore Straits Times Index (SGP), the UK's FTSE 100 (UK), and the US S&P 500 (USA) from 1 April 2008 to 29 July 2011 ($n = 733$ observations).

Univariate time series models

Using the inference functions for margins method outlined in Section 3.2, we modelled univariate marginal distributions without regard to the dependence structure between variables. We first removed serial correlation in the ten time series of log-returns by applying an AR(1)-GARCH(1,1) filter, which accounts for conditional heteroskedasticity present in the data, see Bollerslev (1986). The log-return $r_{i,t}$ of stock index $i \in \{\mathrm{AUS}, \mathrm{CAN}, \mathrm{CH}, \mathrm{DEU}, \mathrm{FRA}, \mathrm{HK}, \mathrm{JPN}, \mathrm{SGP}, \mathrm{UK}, \mathrm{USA}\}$ at time t can thus be written as

$$r_{i,t} = \mu_i + a_i\, r_{i,t-1} + \varepsilon_{i,t}, \quad \varepsilon_{i,t} = \sigma_{i,t}\, z_{i,t}, \quad \sigma_{i,t}^2 = \omega_i + \alpha_i\, \varepsilon_{i,t-1}^2 + \beta_i\, \sigma_{i,t-1}^2,$$

with parameters $\omega_i > 0$, $\alpha_i, \beta_i \geq 0$ such that $\alpha_i + \beta_i < 1$, $|a_i| < 1$, and $\mu_i \in \mathbb{R}$, where $\mathbb{E}\left[z_{t,i}\right] = 0$ and $\mathbb{V}\mathrm{ar}\left[z_{t,i}\right] = 1$. The standardised residuals $z_{i,t}$ are assumed to follow a skewed Student's t distribution with ν_i degrees of freedom and skewness parameter γ_i, see McNeil et al. (2005, Section 3.2). The corresponding cdf will be denoted by $\mathrm{t}_{\nu_i,\gamma_i}$. ML parameter estimates and corresponding standard errors derived from numerical evaluation of the Hessian of the AR(1)-GARCH(1,1) parameters are given in Appendix A. We assessed model fit using the following statistical tests: the Ljung-Box test (Ljung and Box, 1978) with null hypothesis that there is no autocorrelation left in the residuals and squared residuals, the Langrange-multiplier ARCH test (Engle, 1982) with null hypothesis that the residuals exhibit no conditional heteroskedasticity, and the Kolmogorov-Smirnov test (Conover, 1999, Section 6.2) with null hypothesis that the residuals follow a skewed Student's t distribution. None of these null hypotheses could be rejected at the 5% significance level. We then transformed the standardised residuals to uniformly distributed observations $u_{i,t} := \mathrm{t}_{\nu_i,\gamma_i}\left(\sqrt{\frac{\nu_i}{\nu_i-2} + \frac{2\,\nu_i^2\,\gamma_i^2}{(\nu_i-2)^2\,(\nu_i-4)}}\, z_{i,t}\right)$, before modelling the joint dependence structure of the ten time series of log-returns by a PCBN.

Estimating the conditional independence structure with the PC algorithm

We estimated the conditional independence structure of the ten time series of log-returns by applying the PC algorithm with either of the ten conditional independence tests COR, C-GR,

C-H, C-K, D-GR, D-H, D-K, R-GR, R-H, and R-K described in Section 5 (with notation as in Section 6) to the transformed observations $u_{i,t}$. All tests were performed at the 5% significance level. As a result, we obtained three different essential graphs $\mathcal{D}^e_{\mathrm{COR}}$, $\mathcal{D}^e_{\mathrm{GR,H}}$, and $\mathcal{D}^e_{\mathrm{K}}$, of which the first was returned by the PC algorithm with COR, the second was returned by the PC algorithm with either of C-GR, C-H, D-GR, D-H, R-GR, and R-H, and the third was returned by the PC algorithm with either of C-K, D-K, and R-K, respectively. Obviously, a restriction of the class of R-vines to C- or D-vines had not influence on the resulting essential graph. We then oriented undirected edges in the obtained essential graphs, as described in Section 2, in order to obtain DAGs $\mathcal{D}_{\mathrm{COR}}$, $\mathcal{D}_{\mathrm{GR,H}}$, and $\mathcal{D}_{\mathrm{K}}$ from the Markov-equivalence classes represented by $\mathcal{D}^e_{\mathrm{COR}}$, $\mathcal{D}^e_{\mathrm{GR,H}}$, and $\mathcal{D}^e_{\mathrm{K}}$, respectively. More precisely, $\mathcal{D}^e_{\mathrm{COR}}$ contained the two undirected edges AUS $-$ HK and CH $-$ DEU, which we replaced by AUS $\rightarrow$ HK and CH $\rightarrow$ DEU, respectively, based on the heuristic rule that $\mathcal{D}^e_{\mathrm{GR,H}}$ and $\mathcal{D}^e_{\mathrm{K}}$ already contained AUS $\rightarrow$ HK and CH $\rightarrow$ DEU. Similarly, we oriented AUS $-$ JPN into AUS $\leftarrow$ JPN in $\mathcal{D}_{\mathrm{GR,H}}$ and $\mathcal{D}_{\mathrm{K}}$ since $\mathcal{D}^e_{\mathrm{COR}}$ already contained AUS $\leftarrow$ JPN. The DAGs $\mathcal{D}_{\mathrm{COR}}$, $\mathcal{D}_{\mathrm{GR,H}}$, and $\mathcal{D}_{\mathrm{K}}$ are given in Figure 8.

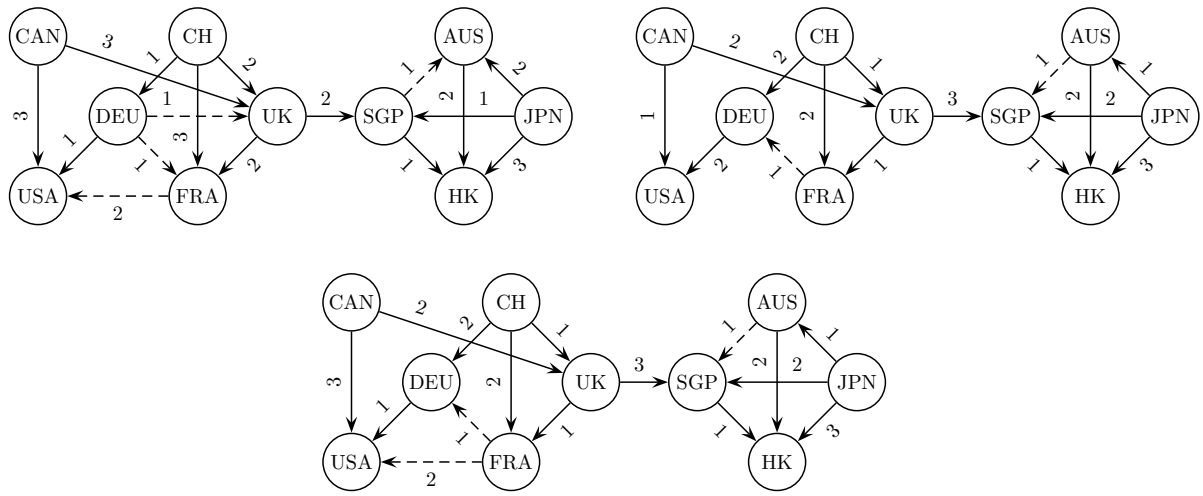

Figure 8: DAGs $\mathcal{D}_{\mathrm{COR}}$ (top left), $\mathcal{D}_{\mathrm{GR,H}}$ (top right), $\mathcal{D}_{\mathrm{K}}$ (bottom) returned by the PC algorithm with different conditional independence tests when estimating the Markov structure of the ten time series AUS, CAN, CH, DEU, FRA, HK, JPN, SGP, UK, USA of daily log-returns. Solid edges appear in all three DAGs. Edge labels indicate parent orderings, that is, for instance, CAN $<_{\mathrm{USA}}$ DEU $<_{\mathrm{USA}}$ FRA in $\mathcal{D}_{\mathrm{COR}}$.

In all three DAGs in Figure 8, the Asian-Pacific indices AUS, HK, JPN, and SGP are mutually adjacent, and so are the two North American indices CAN and USA. The same holds true for the European indices CH, DEU, FRA, and UK in DAG $\mathcal{D}_{\mathrm{COR}}$, while DEU and UK are non-adjacent in $\mathcal{D}_{\mathrm{GR,H}}$ and $\mathcal{D}_{\mathrm{K}}$. A probability measure satisfying the Markov properties represented by either $\mathcal{D}_{\mathrm{GR,H}}$ or $\mathcal{D}_{\mathrm{K}}$, respectively, observes the conditional independence restriction DEU $\perp\!\!\!\perp$ UK $\mid \{\mathrm{CH}, \mathrm{FRA}\}$. All further conditional independence restrictions represented by the DAGs in Figure 8 involve indices in at least two of the above given regions Asia-Pacific, Europe, and North

America. We hence observe a strong geographical clustering of dependences. Moreover, all three DAGs in Figure 8 represent the conditional independence restriction $\{\mathrm{AUS}, \mathrm{HK}, \mathrm{JPN}, \mathrm{SGP}\} \perp\!\!\!\perp \{\mathrm{CAN}, \mathrm{USA}\} \mid \{\mathrm{CH}, \mathrm{DEU}, \mathrm{FRA}, \mathrm{UK}\}$, that is, Asia-Pacific $\perp\!\!\!\perp$ North America $\mid$ Europe. Note that Markov properties alone are not sufficient for deriving causal relations within the analysed data (see, for instance, the undirected edges in an essential graph), but they can be used as a starting point for further research in that direction.

A well-ordering for $\mathcal{D}_{\mathrm{COR}}$ is given by $1 \mapsto$ CAN, $2 \mapsto$ CH, $3 \mapsto$ DEU, $4 \mapsto$ UK, $5 \mapsto$ FRA, $6 \mapsto$ USA, $7 \mapsto$ JPN, $8 \mapsto$ SGP, $9 \mapsto$ AUS, $10 \mapsto$ HK. Similarly, we obtain a well-ordering for $\mathcal{D}_{\mathrm{GR,H}}$ and $\mathcal{D}_{\mathrm{K}}$, respectively, by mapping $1 \mapsto$ CAN, $2 \mapsto$ CH, $3 \mapsto$ UK, $4 \mapsto$ FRA, $5 \mapsto$ DEU, $6 \mapsto$ USA, $7 \mapsto$ JPN, $8 \mapsto$ AUS, $9 \mapsto$ SGP, $10 \mapsto$ HK. We determined parent orderings for the three DAGs in Figure 8 in two steps. First, we applied the greedy-type procedure with Kendall's τ edge weights described in Section 4.2, and second, we permuted some of the orderings obtained in step one to reduce the number of integrals in the corresponding pair-copula decompositions and thus the computational complexity. More precisely, we changed JPN $<_{\mathrm{AUS}}$ SGP and CAN $<_{\mathrm{USA}}$ DEU $<_{\mathrm{USA}}$ FRA in DAG $\mathcal{D}_{\mathrm{COR}}$ into SGP $<_{\mathrm{AUS}}$ JPN and DEU $<_{\mathrm{USA}}$ FRA $<_{\mathrm{USA}}$ CAN, respectively, and CAN $<_{\mathrm{USA}}$ DEU $<_{\mathrm{USA}}$ FRA in DAG $\mathcal{D}_{\mathrm{K}}$ into DEU $<_{\mathrm{USA}}$ FRA $<_{\mathrm{USA}}$ CAN. The resulting parent orderings for $\mathcal{D}_{\mathrm{COR}}$, $\mathcal{D}_{\mathrm{GR,H}}$, and $\mathcal{D}_{\mathrm{K}}$, respectively, are displayed in Figure 8.

Pair-copula selection and ML estimation

Having fixed the parent orderings for the three PCBNs corresponding to $\mathcal{D}_{\mathrm{COR}}$, $\mathcal{D}_{\mathrm{GR,H}}$, and $\mathcal{D}_{\mathrm{K}}$, respectively, we next selected parametric copula families using the AIC as a selection criterion. We considered the Clayton, Frank, Gaussian, Gumbel, and Student's t copula families as well as reflected versions of the Clayton and Gumbel copula families in order to account for negative correlations. We then computed sequential ML estimates of the parameters of the so specified PCBNs. Selected pair-copula families, corresponding sequential ML estimates, bootstrapped standard errors, and estimates of Kendall's τ are given in Table 5. The respective maximised log-likelihoods and AIC values are summarised in Table 6. Moreover, we compared model fit to the respective Gaussian PCBNs comprising only Gaussian pair copulas. Corresponding ML estimates, standard errors, and estimates of Kendall's τ are again found in Table 5, while maximised log-likelihoods and AIC values are given in Table 6.

According to the AIC, the best fit was obtained by the non-Gaussian PCBN with DAG $\mathcal{D}_{\mathrm{GR,H}}$, followed by the non-Gaussian PCBNs associated to $\mathcal{D}_{\mathrm{K}}$ and $\mathcal{D}_{\mathrm{COR}}$, respectively. Applying the Vuong test with AIC correction (Vuong, 1989) to the non-Gaussian PCBNs at the 5% level, we cannot reject the null hypothesis that all three models are equally close to the true model. A similar statement holds for the Gaussian PCBNs. However, using the Vuong test for model

DAG		$\mathcal{D}_{\mathrm{COR}}$			$\mathcal{D}_{\mathrm{GR,H}}$			$\mathcal{D}_{\mathrm{K}}$		
			Parameters	$\widehat{\tau}$		Parameters	$\widehat{\tau}$		Parameters	$\widehat{\tau}$
JPN → AUS	n$\mathcal{G}$	t	0.56, 10.3 (0.03, 4.1)	0.38	t	0.73, 8.5 (0.02, 3.4)	0.52	t	0.73, 8.5 (0.02, 3.4)	0.52
	$\mathcal{G}$	N	0.55 (0.03)	0.37	N	0.72 (0.02)	0.52	N	0.72 (0.02)	0.52
SGP → AUS	n$\mathcal{G}$	t	0.64, 9.4 (0.02, 4.1)	0.44						
	$\mathcal{G}$	N	0.64 (0.02)	0.44						
CH → DEU	n$\mathcal{G}$	t	0.83, 5.4 (0.01, 1.8)	0.63	F	0.98 (0.23)	0.11	F	0.98 (0.23)	0.11
	$\mathcal{G}$	N	0.82 (0.01)	0.62	N	0.14 (0.04)	0.09	N	0.14 (0.04)	0.09
FRA → DEU	n$\mathcal{G}$				t	0.94, 3.8 (0.01, 1.2)	0.78	t	0.94, 3.8 (0.01, 1.2)	0.78
	$\mathcal{G}$				N	0.93 (0.01)	0.76	N	0.93 (0.01)	0.76
CH → FRA	n$\mathcal{G}$	F	1.55 (0.24)	0.17	t	0.44, 9.0 (0.03, 4.4)	0.29	t	0.44, 9.0 (0.03, 4.3)	0.29
	$\mathcal{G}$	N	0.28 (0.04)	0.18	N	0.44 (0.03)	0.29	N	0.44 (0.03)	0.29
DEU → FRA	n$\mathcal{G}$	t	0.94, 3.8 (0.01, 1.1)	0.78						
	$\mathcal{G}$	N	0.93 (0.01)	0.76						
UK → FRA	n$\mathcal{G}$	t	0.57, 6.9 (0.03, 3.3)	0.38	t	0.92, 7.3 (0.01, 2.8)	0.74	t	0.92, 7.3 (0.01, 2.8)	0.74
	$\mathcal{G}$	N	0.59 (0.04)	0.40	N	0.92 (0.01)	0.74	N	0.92 (0.01)	0.74
AUS → HK	n$\mathcal{G}$	t	0.35, 13.4 (0.03, 4.8)	0.22	t	0.35, 13.4 (0.03, 4.8)	0.22	t	0.35, 13.4 (0.03, 4.8)	0.22
	$\mathcal{G}$	N	0.36 (0.04)	0.24	N	0.36 (0.04)	0.24	N	0.36 (0.04)	0.24
JPN → HK	n$\mathcal{G}$	t	0.20, 20.0 (0.04, 2.8)	0.13	N	0.20 (0.04)	0.13	N	0.20 (0.04)	0.13
	$\mathcal{G}$	N	0.20 (0.04)	0.13	N	0.20 (0.04)	0.13	N	0.20 (0.04)	0.13
SGP → HK	n$\mathcal{G}$	t	0.78, 5.7 (0.01, 1.9)	0.57	t	0.78, 5.7 (0.01, 2.0)	0.57	t	0.78, 5.7 (0.01, 2.0)	0.57
	$\mathcal{G}$	N	0.78 (0.02)	0.57	N	0.78 (0.02)	0.57	N	0.78 (0.02)	0.57
AUS → SGP	n$\mathcal{G}$				t	0.64, 9.4 (0.02, 4.0)	0.44	t	0.64, 9.4 (0.02, 4.0)	0.44
	$\mathcal{G}$				N	0.64 (0.02)	0.44	N	0.64 (0.02)	0.44
JPN → SGP	n$\mathcal{G}$	t	0.60, 11.6 (0.02, 4.5)	0.41	N	0.27 (0.03)	0.17	N	0.27 (0.03)	0.17
	$\mathcal{G}$	N	0.60 (0.02)	0.41	N	0.26 (0.03)	0.17	N	0.26 (0.04)	0.17
UK → SGP	n$\mathcal{G}$	N	0.34 (0.03)	0.22	N	0.27 (0.03)	0.18	N	0.27 (0.03)	0.18
	$\mathcal{G}$	N	0.34 (0.03)	0.22	N	0.27 (0.03)	0.18	N	0.27 (0.03)	0.18
CAN → UK	n$\mathcal{G}$	SG	1.13 (0.03)	0.12	N	0.31 (0.03)	0.20	N	0.31 (0.03)	0.20
	$\mathcal{G}$	N	0.21 (0.04)	0.13	N	0.31 (0.03)	0.20	N	0.31 (0.03)	0.20
CH → UK	n$\mathcal{G}$	t	0.36, 9.6 (0.03, 4.5)	0.23	t	0.83, 8.6 (0.01, 3.7)	0.62	t	0.83, 8.6 (0.01, 3.7)	0.62
	$\mathcal{G}$	N	0.39 (0.04)	0.25	N	0.83 (0.01)	0.62	N	0.83 (0.01)	0.62
DEU → UK	n$\mathcal{G}$	t	0.88, 7.3 (0.01, 3.0)	0.69						
	$\mathcal{G}$	N	0.88 (0.01)	0.68						
CAN → USA	n$\mathcal{G}$	N	0.48 (0.03)	0.32	t	0.75, 6.7 (0.02, 3.0)	0.54	N	0.52 (0.03)	0.35
	$\mathcal{G}$	N	0.48 (0.03)	0.32	N	0.75 (0.01)	0.54	N	0.52 (0.03)	0.35
DEU → USA	n$\mathcal{G}$	t	0.71, 6.8 (0.02, 3.0)	0.51	t	0.47, 12.7 (0.03, 4.9)	0.31	t	0.71, 6.8 (0.02, 3.0)	0.51
	$\mathcal{G}$	N	0.71 (0.02)	0.50	N	0.47 (0.03)	0.31	N	0.71 (0.02)	0.50
FRA → USA	n$\mathcal{G}$	t	0.19, 9.2 (0.04, 4.5)	0.12				t	0.19, 9.2 (0.04, 4.4)	0.12
	$\mathcal{G}$	N	0.21 (0.05)	0.14				N	0.21 (0.05)	0.14

Table 5: Selected pair-copula families, sequential ML estimates, standard errors (parentheses), and estimates of Kendall's τ for the Gaussian ($\mathcal{G}$) and non-Gaussian (n$\mathcal{G}$) PCBNs corresponding to the DAGs in Figure 8. Copulas include the Frank (F), Gaussian (N), Survival-Gumbel (SG), and Student's t (t) pair-copula families.

selection between a Gaussian and a non-Gaussian PCBN will always decide in favor of the non-Gaussian model, which again shows the latter models' superiority. This is, of course, to be expected since financial returns often exhibit heavy-tailed dependence, which is validated here by the low estimates of the degrees-of-freedom parameters of the Student's t copulas.

DAG		LL	# Parameters	AIC
$\mathcal{D}_{\mathrm{COR}}$	n$\mathcal{G}$	3397.0	30	−6734.0
	$\mathcal{G}$	3264.6	17	−6495.3
$\mathcal{D}_{\mathrm{GR,H}}$	n$\mathcal{G}$	3412.8	25	−6775.6
	$\mathcal{G}$	3285.5	15	−6540.9
$\mathcal{D}_{\mathrm{K}}$	n$\mathcal{G}$	3401.9	26	−6751.8
	$\mathcal{G}$	3264.1	16	−6496.3

Table 6: Maximised log-likelihoods, numbers of parameters, and AIC values for the Gaussian ($\mathcal{G}$) and non-Gaussian (n$\mathcal{G}$) PCBNs corresponding to the DAGs in Figure 8. Sequential ML estimates of the corresponding parameters are given in Table 5.

8. Conclusion

We have investigated a novel procedure for constructing non-Gaussian continuous Bayesian networks that uses bivariate copulas as building blocks. The resulting models can accommodate a great variety of distributional features to be modelled such as tail-dependence and non-linear, asymmetric dependence. We have provided an algorithm for deriving explicit representations of the corresponding log-likelihoods, as well as routines for random sampling and model selection.

Depending on the underlying DAG and the corresponding parent orderings, the evaluation of the log-likelihood of a PCBN may involve high-dimensional numerical integration and hence considerable computational effort. We have presented a greedy procedure for selecting the parent orderings of the vertices of the underlying DAG, which is based on the idea of modelling strongest dependences in the unconditional pair-copulas. In Section 7, we introduced an additional selection step, in which some of the parent sets were rearranged in order to reduce the number of integrals in the corresponding likelihood decompositions. It would be desirable to have theoretical results on the relationship between parent orderings and the number and complexity of integrals. Bauer et al. (2012) suggested to replace some or all of the integrals by non-parametric kernel conditional cdf estimators. Another way of reducing computational complexity is to consider sequential instead of joint ML estimates.

We used vine copula models to derive a novel test for conditional independence of continuous random variables. The quality of the test, by design, greatly benefits from the ongoing research on vine copulas. In combination with the PC algorithm, we obtained a structure estimation procedure for non-Gaussian PCBNs, which proved to be reliable in the simulation study in Section 6. One may investigate the performance of other conditional independence tests like Zhang et al. (2011), as well as of other structure estimation algorithms. Also, recall that by Meek (1995), constraint-based estimation algorithms can be adapted to incorporate existing expert knowledge. The distributional flexibility of pair-copula Bayesian networks may become even more apparent in application areas other than finance.

Acknowledgements

The authors are very grateful to Peter Hepperger for his help in implementing the algorithms of Section 4 in `C++`. The computer programs were tested on a Linux cluster supported by the DFG (German Research Foundation). Alexander Bauer acknowledges the support of the TUM Graduate School's Faculty Graduate Center ISAM (International School of Applied Mathematics) at the Technische Universität München.

References

K. Aas, C. Czado, A. Frigessi, and H. Bakken. Pair-copula constructions of multiple dependence. *Insurance: Mathematics and Economics*, 44:182–198, 2009.

E. F. Acar, C. Genest, and J. Nešlehová. Beyond simplified pair-copula constructions. *Journal of Multivariate Analysis*, 110:74–90, 2012.

T. W. Anderson. *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons, Chichester, third edition, 2003.

S. A. Andersson, D. Madigan, and M. D. Perlman. A characterization of Markov equivalence classes for acyclic digraphs. *The Annals of Statistics*, 25(2):505–541, 1997.

D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook. *The Traveling Salesman Problem: A Computational Study*. Princeton University Press, Princeton, New Jersey, 2007.

A. Bauer, C. Czado, and T. Klein. Pair-copula constructions for non-Gaussian DAG models. *The Canadian Journal of Statistics*, 40(1):86–109, 2012.

T. Bedford and R. M. Cooke. Probability density decomposition for conditionally dependent random variables modeled by vines. *Annals of Mathematics and Artificial Intelligence*, 32: 245–268, 2001.

T. Bedford and R. M. Cooke. Vines—A new graphical model for dependent random variables. *The Annals of Statistics*, 30(4):1031–1068, 2002.

W. P. Bergsma. Nonparametric testing of conditional independence by means of the partial copula. arXiv:1101.4607v1 [math.ST], 2011.

J. R. Blum, J. Kiefer, and M. Rosenblatt. Distribution free tests of independence based on the sample distribution function. *The Annals of Mathematical Statistics*, 32(2):485–498, 1961.

T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.

E. C. Brechmann, C. Czado, and K. Aas. Truncated regular vines in high dimensions with application to financial data. *The Canadian Journal of Statistics*, 40(1):68–85, 2012.

W. J. Conover. *Practical Nonparametric Statistics*. John Wiley & Sons, New York, third edition, 1999.

R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems: Exact Computational Methods for Bayesian Networks*. Springer, New York, second edition, 2003.

C. Czado, U. Schepsmeier, and A. Min. Maximum likelihood estimation of mixed C-vines with application to exchange rates. *Statistical Modelling*, 12(3):229–255, 2012.

J. Dißmann, E. C. Brechmann, C. Czado, and D. Kurowicka. Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics and Data Analysis*, 2012. To appear.

D. Dor and M. Tarsi. A simple algorithm to construct a consistent extension of a partially oriented graph. Technical report R-185, Cognitive Systems Laboratory, UCLA, 1992.

G. Elidan. Copula Bayesian networks. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta (Eds.), *Advances in Neural Information Processing Systems*, volume 23, pages 559–567. NIPS Foundation, La Jolla, California, 2010.

G. Elidan. Lightning-speed structure learning of nonlinear continuous networks. In N. Lawrence and M. Girolami (Eds.), *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *JMLR: Workshop and Conference Proceedings*, pages 355–363. Journal of Machine Learning Research, 2012.

R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

C. Genest and B. Rémillard. Tests of independence and randomness based on the empirical copula process. *Test*, 13(2):335–369, 2004.

C. Genest, K. Ghoudi, and L.-P. Rivest. A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3):543–552, 1995.

L. Gruber, C. Czado, and J. Stöber. Bayesian model selection for regular vine copulas using reversible jump MCMC. Submitted for publication, 2012.

A. M. Hanea and D. Kurowicka. Mixed non-parametric continuous and discrete Bayesian belief nets. In T. Bedford, J. Quigley, L. Walls, B. Alkali, A. Daneshkhah, and G. Hardman (Eds.), *Advances in Mathematical Modeling for Reliability*, pages 9–16. IOS Press, Amsterdam, 2008.

A. M. Hanea, D. Kurowicka, and R. M. Cooke. Hybrid method for quantifying and analyzing Bayesian belief nets. *Quality and Reliability Engineering International*, 22:709–729, 2006.

A. M. Hanea, D. Kurowicka, R. M. Cooke, and D. A. Ababei. Mining and visualising ordinal data with non-parametric continuous BBNs. *Computational Statistics and Data Analysis*, 54 (3):668–687, 2010.

N. Harris and M. Drton. PC algorithm for Gaussian copula graphical models. arXiv:1207.0242v1 [math.ST], 2012.

I. Hobæk Haff. Parameter estimation for pair-copula constructions. *Bernoulli*, 2012a. To appear.

I. Hobæk Haff. Comparing estimators for pair-copula constructions. *Journal of Multivariate Analysis*, 110:91–105, 2012b.

I. Hobæk Haff, K. Aas, and A. Frigessi. On the simplified pair-copula construction—Simply useful or too simplistic? *Journal of Multivariate Analysis*, 101(5):1296–1310, 2010.

W. Hoeffding. A non-parametric test of independence. *The Annals of Mathematical Statistics*, 19(4):546–557, 1948.

M. Hollander and D. A. Wolfe. *Nonparametric Statistical Methods*. John Wiley & Sons, Chichester, second edition, 1999.

H. Joe. Families of m-variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters. In L. Rüschendorf, B. Schweizer, and M. D. Taylor (Eds.), *Distributions with Fixed Marginals and Related Topics*, volume 28 of *Lecture Notes—Monograph Series*, pages 120–141. Institute of Mathematical Statistics, Hayward, California, 1996.

H. Joe. *Multivariate Models and Dependence Concepts*. Chapman & Hall, London, 1997.

H. Joe and J. J. Xu. The estimation method of inference functions for margins for multivariate models. Technical report 166, Department of Statistics, University of British Columbia, 1996.

M. Kalisch and P. Bühlmann. Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *Journal of Machine Learning Research*, 8:613–636, 2007.

M. Kalisch, M. Mächler, and D. Colombo. *Estimation of CPDAG/PAG and causal inference using the IDA algorithm*, 2012. Manual for the R package `pcalg`.

G. Kim, M. J. Silvapulle, and P. Silvapulle. Comparison of semiparametric and parametric methods for estimating copulas. *Computational Statistics and Data Analysis*, 51(6):2836–2850, 2007.

D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques.* MIT Press, Cambridge, Massachusetts, 2009.

J. B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society*, 7(1):48–50, 1956.

D. Kurowicka. Optimal Truncation of Vines. In D. Kurowicka and H. Joe (Eds.), *Dependence Modeling: Vine Copula Handbook*, pages 233–248. World Scientific, Singapore, 2011.

D. Kurowicka and R. M. Cooke. Distribution-free continuous Bayesian belief nets. In A. Wilson, N. Limnios, S. Keller-McNulty, and Y. Armijo (Eds.), *Modern Statistical and Mathematical Methods in Reliability*, volume 10 of *Series on Quality, Reliability and Engineering Statistics*, pages 309–323. World Scientific, Singapore, 2005.

D. Kurowicka and H. Joe (Eds.). *Dependence Modeling: Vine Copula Handbook.* World Scientific, Singapore, 2011.

S. L. Lauritzen. *Graphical Models.* Oxford University Press, Oxford, 1996.

G. M. Ljung and G. E. P. Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978.

A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools.* Princeton University Press, Princeton, New Jersey, 2005.

C. Meek. Causal inference and causal explanation with background knowledge. In P. Besnard and S. Hanks (Eds.), *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pages 403–410. Morgan Kaufmann, San Francisco, California, 1995.

A. Min and C. Czado. Bayesian model selection for D-vine pair-copula constructions. *The Canadian Journal of Statistics*, 39(2):239–258, 2011.

B. R. Moole and M. Valtorta. Sequential and parallel algorithms for causal explanation with background knowledge. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 12(Supplement 2):101–122, 2004.

R. E. Neapolitan. *Learning Bayesian Networks.* Prentice Hall, Upper Saddle River, New Jersey, 2003.

R. B. Nelsen. *An Introduction to Copulas.* Springer, New York, second edition, 2006.

J. Pearl. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, Cambridge, second edition, 2009.

O. Pourret, P. Naïm, and B. Marcot (Eds.). *Bayesian Networks: A Practical Guide to Applications*. John Wiley & Sons, Chichester, 2008.

R. C. Prim. Shortest connection networks and some generalizations. *Bell System Technical Journal*, 36(6):1389–1401, 1957.

R. W. Robinson. Counting labeled acyclic digraphs. In F. Harary (Ed.), *New Directions in the Theory of Graphs*, pages 239–273. Academic Press, New York, 1973.

M. Rosenblatt. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*, 23(3):470–472, 1952.

U. Schepsmeier, J. Stöber, and E. C. Brechmann. *Statistical inference of vine copulas*, 2012. Manual for the R package `VineCopula`.

A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231, 1959.

M. Smith, A. Min, C. Almeida, and C. Czado. Modeling longitudinal data using a pair-copula decomposition of serial dependence. *Journal of the American Statistical Association*, 105(492): 1467–1479, 2010.

K. Song. Testing conditional independence via Rosenblatt transforms. *The Annals of Statistics*, 37(6B):4011–4045, 2009.

P. Spirtes and C. Glymour. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9(1):62–72, 1991.

P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT Press, Cambridge, Massachusetts, second edition, 2000.

I. Tsamardinos, L. E. Brown, and C. F. Aliferis. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, 65(1):31–78, 2006.

T. Verma and J. Pearl. Equivalence and synthesis of causal models. In P. Bonissone, M. Henrion, L. Kanal, and J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence*, volume 6, pages 220–227. Elsevier Science, New York, 1991.

Q. H. Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, 57(2):307–333, 1989.

K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In F. Cozman and A. Pfeffer (Eds.), *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 804–813. AUAI Press, Corvallis, Oregon, 2011.

A. Estimated AR-GARCH parameters and standard errors

	μ [$\times 10^3$]	a	ω [$\times 10^5$]	α	β	ν	γ
AUS	0.26 (0.41)	0.01 (0.04)	0.18 (3.77)	0.08 (0.02)	0.91 (0.02)	10.33 (3.46)	0.90 (0.05)
CAN	0.50 (0.38)	−0.01 (0.04)	0.17 (3.83)	0.09 (0.02)	0.90 (0.02)	10.74 (3.73)	0.80 (0.05)
CH	0.08 (0.38)	0.02 (0.04)	0.38 (4.01)	0.12 (0.03)	0.86 (0.02)	7.64 (2.04)	0.92 (0.05)
DEU	0.66 (0.49)	−0.03 (0.04)	0.31 (4.03)	0.08 (0.02)	0.91 (0.02)	7.77 (2.39)	0.93 (0.05)
FRA	0.23 (0.54)	−0.02 (0.04)	0.59 (4.24)	0.09 (0.02)	0.89 (0.02)	8.29 (2.44)	0.94 (0.05)
HK	0.36 (0.55)	−0.02 (0.04)	0.23 (4.07)	0.07 (0.02)	0.93 (0.01)	8.23 (2.25)	0.97 (0.05)
JPN	0.27 (0.53)	−0.06 (0.04)	0.85 (4.39)	0.12 (0.03)	0.85 (0.02)	15.33 (7.85)	0.87 (0.05)
SGP	0.57 (0.40)	−0.01 (0.04)	0.24 (3.98)	0.09 (0.02)	0.90 (0.02)	4.92 (0.89)	1.03 (0.05)
UK	0.56 (0.44)	−0.01 (0.04)	0.37 (4.02)	0.09 (0.02)	0.89 (0.02)	8.42 (2.60)	0.92 (0.05)
USA	0.76 (0.43)	−0.07 (0.04)	0.23 (4.03)	0.10 (0.02)	0.89 (0.02)	7.22 (2.15)	0.84 (0.04)

Table 7: ML estimates and standard errors (in parentheses) of AR(1)-GARCH(1,1) parameters for the ten time series of daily log-returns analysed in Section 7.